SUMMER 1979 VOL. 1, NO. 1 $2.00 USA/2.40 CANADA

ROBOTICS AGE

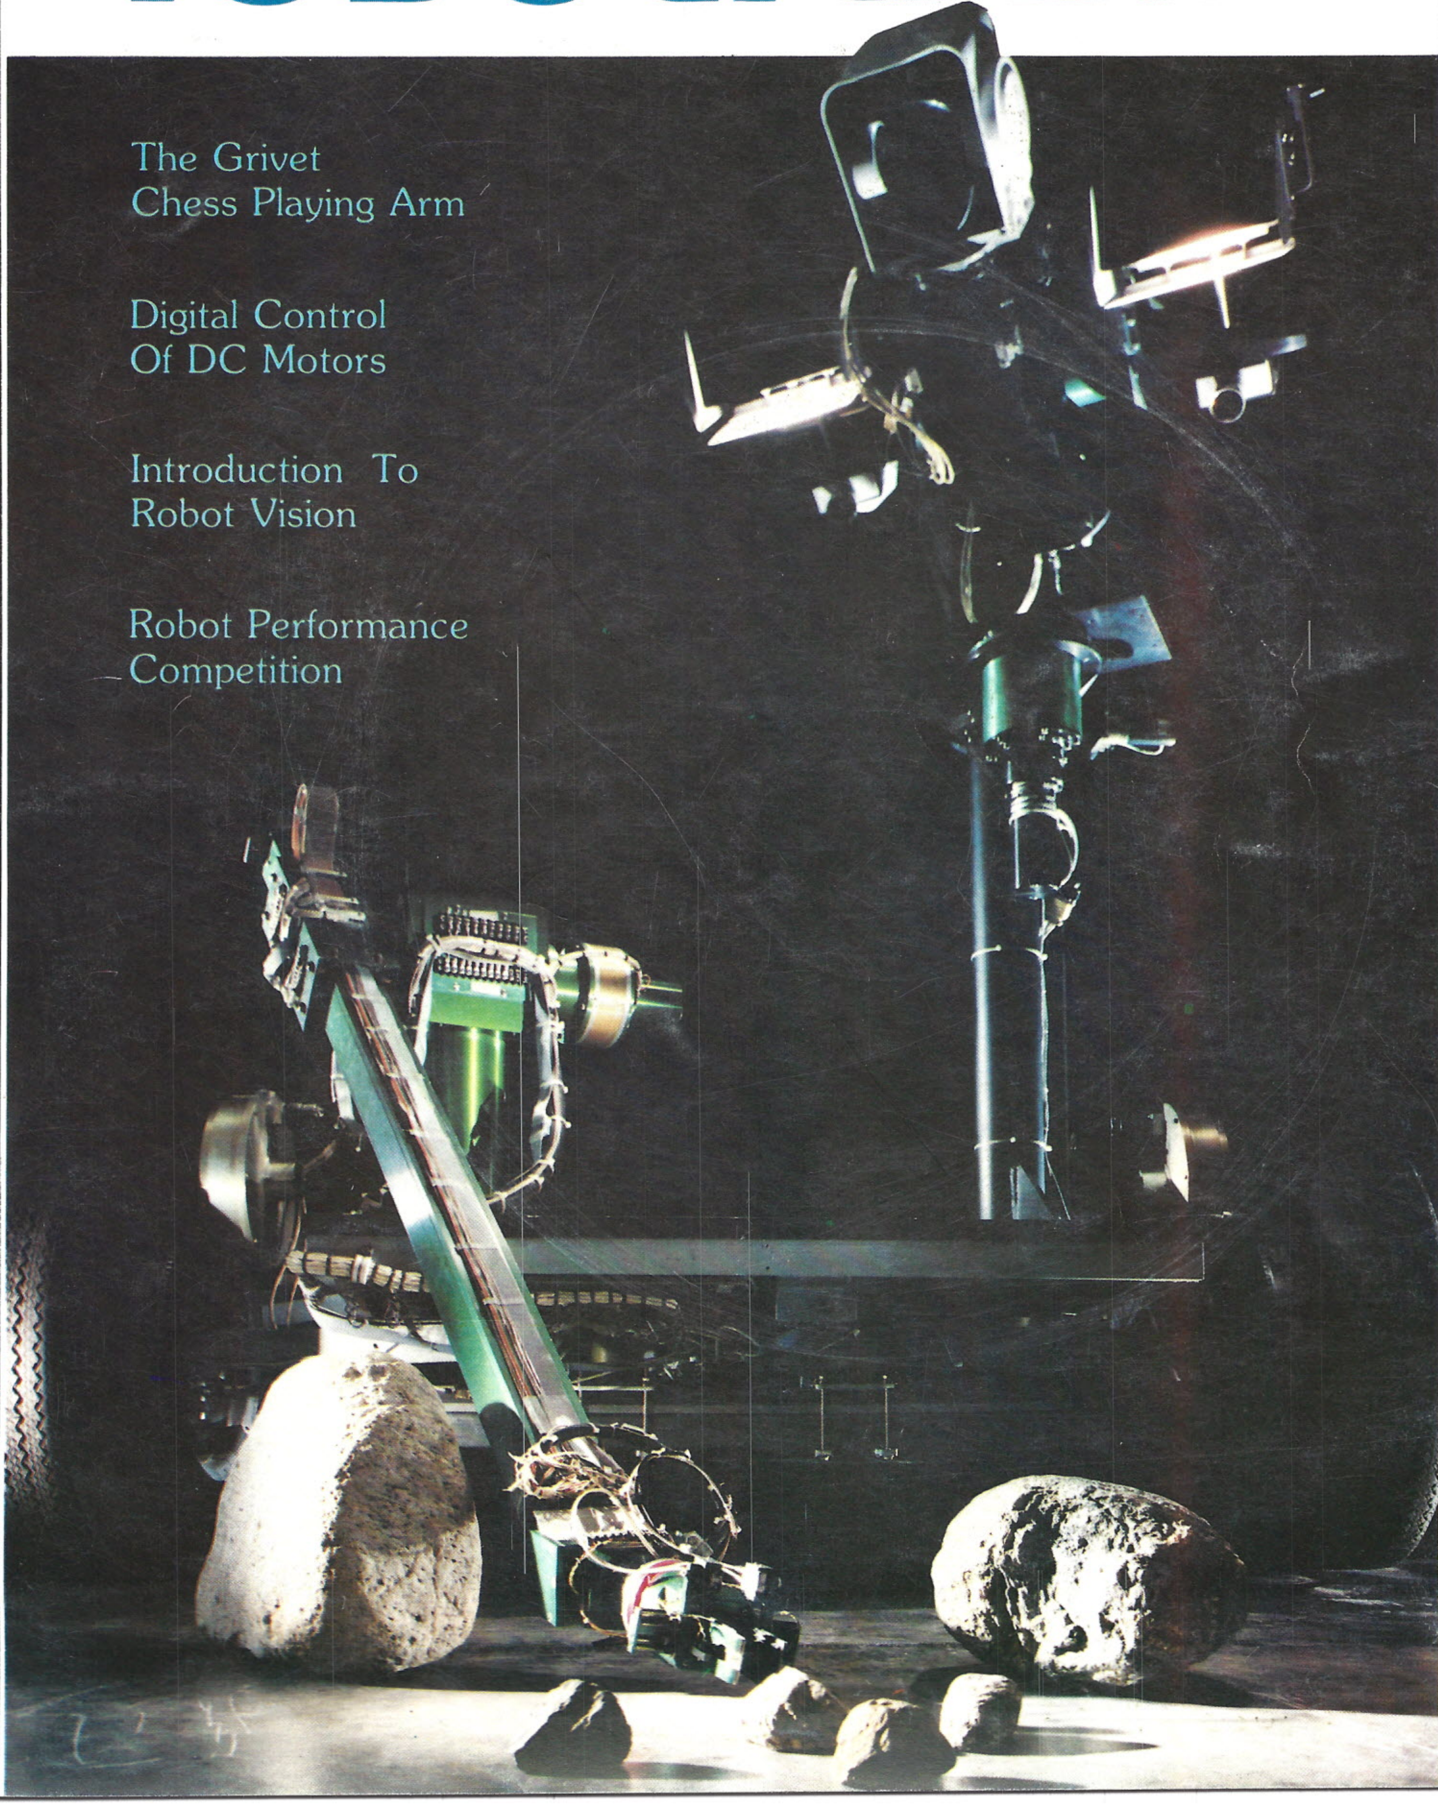

The Grivet
Chess Playing Arm

Digital Control
Of DC Motors

Introduction To
Robot Vision

Robot Performance
Competition

ROBOTICS AGE

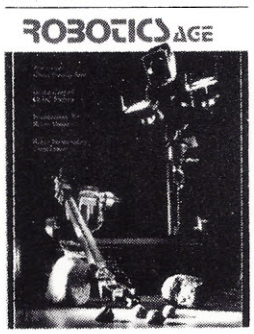

COVER: The research robot of the NASA Jet Propulsion Laboratory Robotics Research Program. This system has two GE solid-state TV cameras, a laser range-finder, and a six-jointed manipulator. The robot can navigate around obstacles and collect visually-located rock samples. It has the ability to visually track several moving targets at once and to pursue one of them. It can also perform simple assembly tasks.
PHOTO COURTESY OF JET PROPULSION LABORATORY

Articles

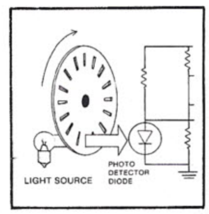

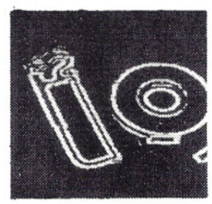

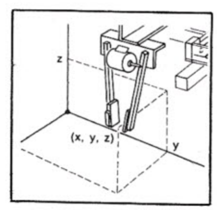

PUBLISHER
Philip C. Flora

EDITORIAL DIRECTOR
Alan M. Thompson

EDITOR
Hugh Bartel

EDITORIAL ASSISTANT
Jeff Daiell

ART DIRECTOR
Robert Tinney

TECHNICAL ART
Andrew Layman

PRODUCTION EDITOR
Cindee Huff

ARTIST
Steve Fagerquist

TYPOGRAPHERS
Robert Bakhaus
Karen Kendall
Craig Haynie

Bruce McDonald

SUMMER 1979 VOL. 1, No. 1

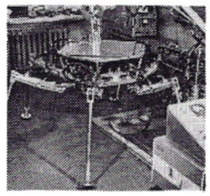

Departments

ROBOTICS AGE is published four times a year, in the summer, fall, winter, and spring by Robotics Publishing Corporation, 3410 Marquart, Suite 203, Houston, Texas 77027. Subscriptions are $7.50 for one year in the USA. Application to mail at Second-Class Postage Rates is pending at Houston, Texas. POSTMASTER: Send address changes to ROBOTICS AGE, P. O. Box 4029, Houston, Texas 77210.

Subscriptions from Canada and Mexico are $10 for one year. Other foreign countries $15 for one year surface delivery (add $16 for airmail delivery). Single copy price is $2.00 in the USA, $2.50 in Canada and Mexico, $4 in Europe, and $4.50 elsewhere. Foreign subscriptions and sales should be remitted in United States funds drawn on a U.S. bank.

Address all correspondence to P. O. Box 4029, Houston, Texas 77210, phone (713) 877-1596. Unacceptable manuscripts will be returned if accompanied by sufficient first class postage. Not responsible for lost manuscripts or photos. Opinions expressed by the authors are not necessarily those of ROBOTICS AGE.

PUBLISHER'S STATEMENT

WELCOME TO THE ROBOTICS AGE

Publishing the first issue of **ROBOTICS AGE** has been one of the most arduous tasks I have ever undertaken. My goal was to put together a magazine that would be useful to everyone with an interest in the field. Professional robotics researchers and robotics industry engineers should find articles on state of the art technology and the newest applications. Amateur experimenters should find articles on new circuits, kits, clubs and competitions. Students should be able to learn about robotics; as should decision makers at firms contemplating capital investment in robot equipment.

In short, researchers and engineers need a central information source, amateurs need to know what others are doing, the robotics industry needs a salesman, and the general public needs to be educated to the vast potentials of the field.

I took this idea to Alan Thompson. Alan, a close friend for many years, is a professional robotics researcher and an excellent writer. He is also responsible (along with Isaac Asimov) for my fascination with robots.

He listened to my proposal and at the end posed only one objection; "The field is growing so fast, if it's not already impossible to cover all that in one publication, it soon will be." I told him, "That's fine with me." We would call the company Robotics Publishing Corporation and accept the challenge of turning this problem into an opportunity. We would spin-off a new publication as often as necessary to satisfy the demands of our readers. He agreed to take on the toughest job of all, Editorial Director, and the "idea" started down the road to becoming a reality.

When Robert Tinney, who paints those superb covers for Byte Magazine, agreed to be our Art Director, we were firmly on course. On this foundation we have built an excellent staff.

The problems have been many, not the least of which was my decision to allow only half the normal lead time and try to get out what should have been a November issue before the end of summer. Instead of a mutiny, everyone, on their own initiative, began to stay late evenings and come in weekends so that **ROBOTICS AGE** could have a Summer '79 issue. I want to take this opportunity to thank them all. And I want to give special thanks to Carl Helmers, Editorial Director of BYTE Magazine, who's advice helped us avoid numerous pitfalls and kept us from having to "re-invent the wheel."

Finally, its interesting to note that Alan was right about the need for spin-offs. Our first, the Robotics Industry Directory, is announced on page 21. And how has the "idea" of a magazine on robotics succeeded? Well, our first issue is sold out and we are doubling the press run on the next.

EDITORIAL

On Misrepresentation, Sensationalism and Robots...

A crowd had gathered in the appliance section of a large Pittsburgh department store, attracted by a demonstration of a "robot" doing housework. Students from the Artificial Intelligence section of the Computer Science Department of Carnegie-Mellon University, a major AI research center, heard of the demo and came to see if this "robot" was capable of performing as claimed. The anthropomorphic machine was busily vacuuming a carpet, working around furniture, and conversing with the amazed spectators.

The students, suspicious of this capability apparently beyond the state of the art in research labs, tried to find out how the "robot" worked. When covering up the machine's "sensors" with a handkerchief made no difference in its behavior, they looked around for signs of remote human control. Looking about the room, they spotted one man over in the corner of the room, a little apart from the crowd, with his hand inside of a bag. Another was spotted elsewhere in the room, his mouth suspiciously concealed by a handkerchief. Closer examination exposed a R/C transmitter in the bag, and a palm mike behind the handkerchief. The "robot" was no more than a remote-controlled toy.

There is certainly a place for teleoperated systems; they are invaluable for work in hazardous environments or inaccessible places. Even the distinction between robots and teleops is a fuzzy one — in many cases, computer assist is used to correct for tracking errors, etc., simplifying the task of direct human control. To be sure, even the most advanced research and industrial robots must respond to commands from humans, so complete autonomy is certainly not a distinguishing requirement for a robot nor even a desirable feature. We are familiar with the numerous systems of classifying "robots" according to their capabilities or appearance, but we feel that the primary feature that distinguishes a robot from a teleop is the system's ability to carry out a task, once commanded, (or programmed) independent of human control.

Obviously there is a demand for teleops masquerading as robots — the special effects for many science fiction movies require such systems, and people want them for entertainment at conventions and parties. There is nothing wrong with this in itself, but many uninformed people are deceived by the effect into believing that the machines are self-contained. We feel that such misunderstanding is unfortunate, especially when media coverage contributes to the sensationalism, as occurred in the Pittsburgh incident. Also, in many cases, misrepresentation of a teleop system offered as a commercial product can constitute fraud.

Our point is that teleop systems, especially those presented as anthropomorphic robots, are not what serious robotics experimentation is all about. Our interests, as will be reflected in the technical content of ROBOTICS AGE, concern the continued development of genuine automated robots, leading to increased capabilities in both sensory processing and computer control. One of our aims is to serve as a central source of technical information on robotics, making the results of current research available to a wide audience. Because of the potential of microcomputer applications to low-budget systems in this field, we feel that in many cases the amateur experimenter may have as much to contribute as the professional researcher, and there is a need to provide introductory articles and discussions of simple systems to help new experimenters get started.

We encourage your contribution of original papers describing innovations in robot design and applications, particularly those that pertain to low-budget hardware and software techniques using microcomputers. We intend to be guided in our selection of articles by the response we get from our readers, so please let us know what kind of articles you want. We want to stimulate and encourage experimentation in robotics at all levels, both because of the obvious beneficial impact it may have on our standard of living, and because of its value as a challenging and exciting avocation.

Alan M Thompson

Inexpensive and effective servo techniques suitable for microcomputer-controlled robots are one of our major interests. This article describes a hardware servo method that is both cheap and highly accurate.

Digital Speed Control of DC Motors

by John Craig
NASA Jet Propulsion Laboratory
Pasadena, CA

Most home robotics hobbyists would agree that immobile robots just aren't fun. Many robots in research labs and in industry are just mechanical arms sticking out of a table or mounted along an assembly line. In contrast, most robots built by hobbyists for fun are capable of locomotion. Although the captive robots may be capable of impressive feats within their limited world, we usually want our robots to wander about the house looking for an outlet to recharge themselves, seeking light, fetching the evening paper, and a host of other tasks, all requiring the ability to move. We may strive to give our robot eyes, ears, or an arm, but the first step is to make it MOVE! Many home robots are barely more than moving boxes, but from that base new features can be added as fast as the inventor can dream them up. The great majority of today's robots that move do so with wheels. Robots that walk on articulated legs are still rare, and are complex, expensive, and, so far, limited in capability. Although nature has found it easier to evolve creatures with legs (for movement on dry land), with today's technology wheels are the best alternative. With the exception of staircases, wheels will allow a robot to rove almost anywhere in the house!

I. Introduction to Feedback Control.

The methods of controlling a motor in a robot can range from the most simple on/off power switch to the complex servo control circuits used in industrial manipulators. If variations in the speed of the motor are not important, but different speed ranges are desired, then a simple approach is just to set the input voltage to the motor with a variable power supply. But suppose you want your robot to roll at the same speed regardless of variations in the load. If you have a fixed voltage setting on the motor, the robot will move faster going down inclines and slower when it is climbing or if it is pushing or pulling a load. More important, if more than one motor is used to drive the wheels, as in the system discussed here, accurate speed control of each motor is critical. When commanded to walk a straight line, each motor must run at the proper speed or the robot will follow an embarassingly crooked path. If the speed of each wheel is accurately known, then a robot with a microcomputer can also "navigate" by keeping track of its position relative to some initial starting point or coordinate frame. This can be done by a numerical integration of the robot's velocity. Put more simply, the formula "distance = rate x time" is used, for which the "rate" or speed must be accurately known during the elapsed "time" of the movement for the calculation to be meaningful.

To insure accurate speed it is essential to "close the loop" by using feedback. The "loop" is the feedback loop, and "closing" it refers to using a "feedback" signal from a

sensor to automatically regulate the input to the drive system. In this case, feedback indicating wheel speed must be used to control the input voltage to the motor. In the general case, however, feedback may be used to control almost any dynamic quality of a system, subject to stability restrictions that can be determined by a mathematical analysis of the problem. A generalized feedback loop is shown in Figure 1. The commanded value of the controlled parameter is set from outside the loop or else could be an intrinsic value of the system, as often occurs in living systems. The commanded value is compared against the measured value of the parameter as determined by the feedback sensor. If there is no difference, then no action need be taken, but if there is, then that difference is "amplified", or otherwise processed, by the "gain" element of the loop to produce the input to the system's effector. The effector is called such because it "affects" reality in some way to accomplish the desired change in the controlled parameter, which is being measured by the sensor.

In complex systems, there may be loops within loops, or measurements from different types of sensors may be used. The gain process may range from simple amplification to involved combination of several feedback terms. For example, suppose it is desired to control the force that a motor-driven effector applies to an object. This might occur if an assembly-line robot has to insert a fragile part in an object under construction, or if you want your personal robot to properly use a sponge! A suitable sensor would be a pressure transducer such as a strain guage, which produces a signal that can be processed to measure the pressure applied by the effector. The difference between the measured pressure and the desired pressure can be integrated, e.g., by an operational amplifier circuit, to determine the current to a DC permanent magnet motor driving the effector (mechanical manipulator, etc.)

Feedback loops of one form or another are among the building blocks of life and are the cornerstones of Robotics. In living systems, feedback loops can be found regulating heartbeat, respiration, blood chemistry, intra-cellular pressure, and many forms of motion. Industrial robots depend on many types of feedback loops, typically position, rate, or force control of manipulators. In advanced robots, visual feedback may be employed to aid in grasping, parts positioning, and other tasks. Feedback loops applied to controlling motors are called servo systems, or servos. For a good introduction to servo design, consult Reference [4] (listed at the end of the article) or one of the many other college texts on control system design.

II. Robot Locomotion Using a Phase Locked Servo

Figure 2 shows a basic system for robot locomotion. The front wheels are motor driven and the two rear wheels are small casters. If the motors are reversible this robot can roll forward, backward, turn right or left, and even do fairly graceful piroettes (one motor driven forward and one backward). Photo 1 shows an example of the two motor drive system. Using a microcomputer to control the robot is not a necessity, but doing so will certainly make any robot more talented through the flexibility and power of software.* Microcomputers and robotics are two hobbies that mix excellently! Single-board microcomputers may now be bought for a few hundred dollars, putting them within the economic reach of most robotics

** Editor's note: As discussed in our Editorial, ROBOTICS AGE adheres to the definition of robots as having internal decision-making and control based on sensory input. By this standard, an uncontrolled or remote-controlled mechanical device doesn't qualify in the strictest sense. However, we want to encourage hobbyists to build their machines, whether they are robots in the formal sense or not; computer control can always be added later. The main goal is to build, learn and enjoy!*

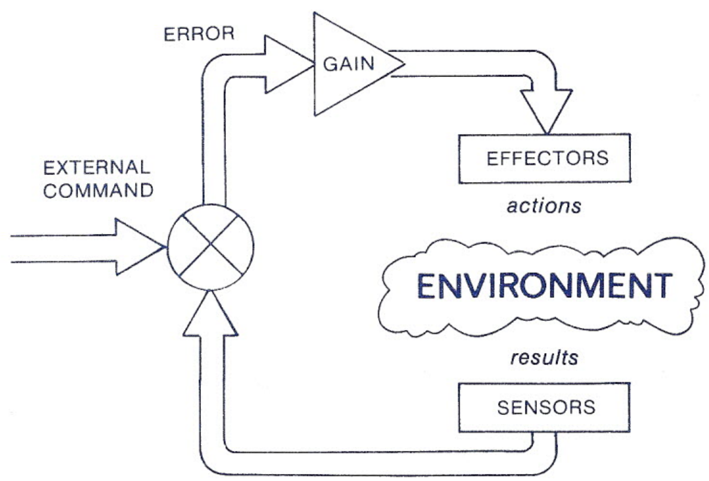

Figure 1. Generalized Feedback Loop: Sensors measure some parameter in the environment. The processed signal is compared with the commanded value, producing an error signal if there is a difference. The error signal is processed by the gain element, resulting in the control signal to the effectors which then alter the environment.

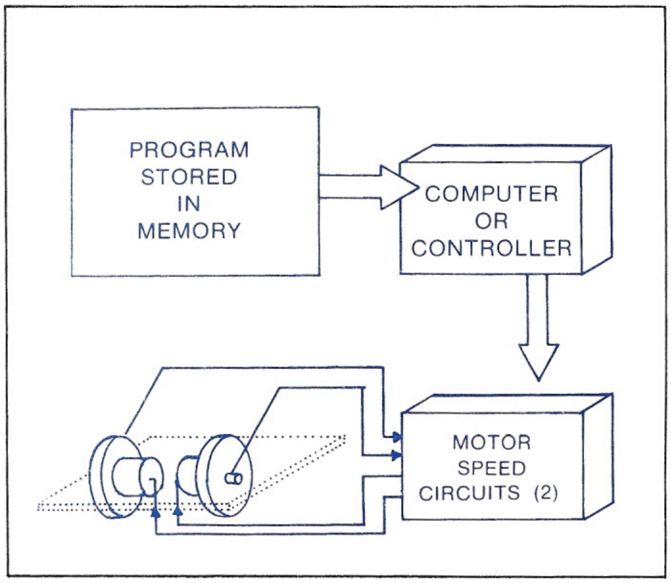

Figure 2. Basic system for robot locomotion: robot is built on a wheeled base which has two powered wheels in front and two free rolling casters in the back. A speed control circuit is needed for each motor to ensure that it turns at the desired speed. A microcomputer can give "high-level" commands to the speed control circuits.

experimenters. Today's microcomputers can, after all, do more than play TV games! The addition of even the smallest microcomputer opens up a whole new world of autonomous operation for your robot when properly interfaced and programmed. But don't be discouraged if your robot doesn't have one, many home built robots haven't had that advantage and still have served their masters well!

The motor speed control circuit is a feedback loop which is given a desired speed command (from an output port of a microcomputer perhaps) and, by comparing it to feedback from the motor, outputs the voltage necessary to drive the motor at that speed. This kind of servo can be realized in many forms. Servos often use feedback from analog tachometers and are designed with analog circuitry which may drift as it ages or with temperature changes and needs to be "tweaked" regularly to be correct. A digital circuit offers many advantages. First, it is easy to connect to a microcomputer, no digital to analog converter is necessary. Digital circuits are much less susceptible to electrical noise than analog circuits, which is an important consideration when working near motors. The accuracy of a digital circuit may be made far superior to that of an analog loop, since it may be driven by a crystal-controlled clock signal.

This article describes a digital motor speed control system called a Phase Locked Servo (PLS). This is a fairly new technique, but the circuitry can be amazingly simple, inexpensive, and extremely accurate. The circuit requires no digital to analog or analog to digital converters. It also can be run off a single voltage power supply, so the need for a multiple-level power supply on board the robot may be avoided. The motor can be made reversible with the addition of a simple relay (mechanical or solid-state). It is a complete hardware servo system; it can be used without a microcomputer, or, if a micro is used, the computer only issues speed commands and is free to do other more interesting tasks while the PLS controls the motor speed.

The only drawback of a PLS system is that it requires a pulse rate tachometer compatible with digital circuitry. Commercially available digital tachs are still fairly expensive (usually over $100). One way around this problem is to build one yourself! An optical encoder is a device which outputs pulses at a rate proportional to the rotational speed of its shaft. It is used as the feedback sensor in the PLS to measure the motor (or wheel) speed. Figure 3 shows a possible way of building your own

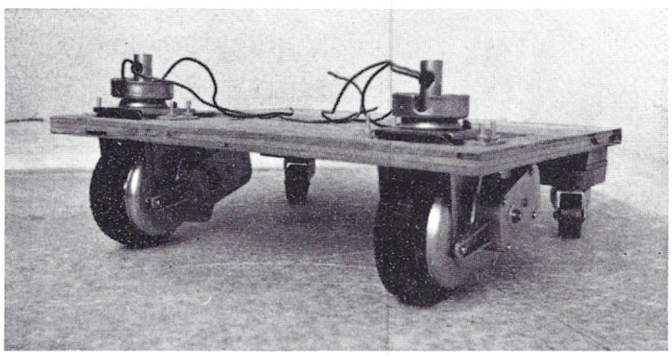

***Photo 1.** Example of the 2-wheel, 2-caster system as implemented on the author's home robot. The powered wheels were bought as units, originally designed for use on electric-powered wheelchairs.*

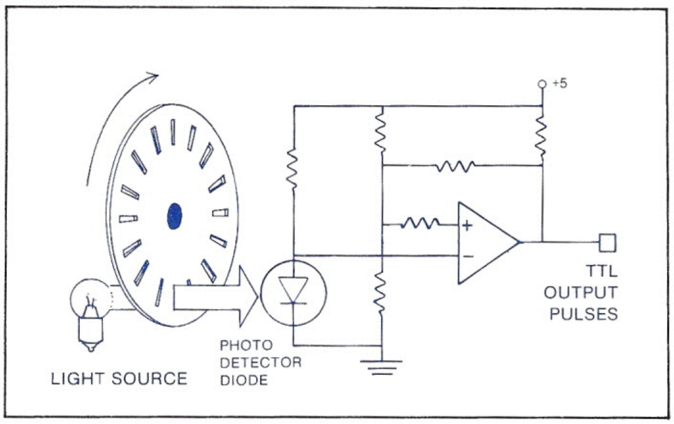

Figure 3. Simple Encoder Circuit: Possible method for converting light pulses to TTL level pulses.

encoder. A disk with evenly spaced holes around its perimeter is mounted on the motor shaft or wheel of the robot. A light source illuminates it from one side so that a photo-detector is illuminated by light pulses as the shaft is turned. The circuit shown is one of many possibilities; it simply converts the light pulses to TTL level digital pulses. You may want to experiment with the circuit to get the best results with your light source and photo-detector.

For a stable system, the number of holes in the disk should be sufficiently large and they should be evenly spaced. If the disk is mounted on a high speed shaft (before a gear reduction to the wheel) the number of holes can be fewer. The number of holes times the disk's rotational speed gives the pulse rate output by the encoder. At the minimum desired shaft speed this value should not be less than about 100 Hz. for good control. If you don't have access to a drill press with a dividing gear to accurately measure the hole spacing, you can geometrically define equally spaced holes by constructing regular polygons on the disk or a template. Start with a perfect square inscribed in your disk and find the midpoint of each chord. The radius through the midpoint of each chord gives the location of four new points on the perimeter, from which eight new chords may be drawn to form an octagon. (See Figure 4.) Continue constructing new chords until the number of perimeter points is the power of two that is near to the number of holes determined by dividing 100 (pulses per second) by your minimum desired shaft rate in revolutions per second. The encoder for the circuit described here has 256 holes. If the holes are too close together for you to drill them accurately, either use a larger disk or put the encoder on a higher speed shaft. Use your construction as a template for drilling or punching uniformly sized holes at a constant radius of your disk (sheet metal, painted plexiglass, etc.) For a given hole spacing, the hole size determines the duty cycle of the output square wave; use small holes with clean edges for a good signal.

Another important element in the PLS is the phase-frequency detector. The one used here is the Motorola MC4044. This circuit accepts two input waveforms and generates an error voltage that is proportional to the frequency and/or phase difference of the input signals. The error voltage is zero when the feedback signal matches the reference signal in phase and frequency but grows as the feedback deviates from the reference. If only phase difference alone was measured, the system could lock onto multiples of the reference frequency. Using the MC4044, this cannot occur, so that the motor may be started from zero and will always reach the right speed, provided, that is, that your power supply is sufficient. The MC4044 also contains an amplifier which will serve in the gain element of the PLS system.

Figure 5 shows a block diagram of a PLS system. The speed command is a 3 bit binary number. The first bit specifies the motor's direction and the other 2 bits select one of four speeds. A clock circuit generates a reference signal which is divided down appropriately in accordance with the command. The resulting square wave is used as the reference input to the phase-frequency detector. The frequency at which the master clock runs must be selected by/from the encoder's characteristics and the desired speed of the motor. If the encoder produces N pulses per revolution, and the maximum (encoder) shaft speed desired is R revolutions per second, then the clock frequency should be N*R cycles per second. By dividing this frequency down, the lower reference frequencies are obtained. In the circuit shown, the minimum shaft speed (used above to determine the proper number of holes in the encoder disk) is one fourth the maximum speed. The reference signal is compared with the output of the encoder, generating an error signal at the output of the phase-frequency detector. The error signal must be filtered to remove the high frequency components. Its average DC level is used as the input to an inverting amplifier which drives the motor. The system is self-regulating; when the error signal grows in magnitude, the motor's voltage is raised or lowered to reduce the error. If the circuit is operating properly there will be a small offset error in phase, giving an average DC level to the amplifier, but almost no error in velocity. PLS systems can achieve

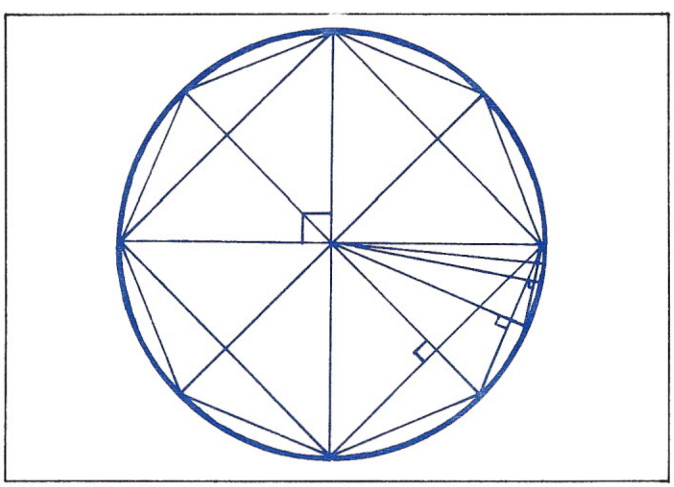

Figure 4. A method of constructing regular polygons of 2^n sides. Used to locate positions of holes in encoder disk.

speed accuracies as good as 0.02%!

A working circuit is shown in Figure 6. The master clock is implemented using the familiar NE555 timer integrated circuit. The values of R_a, R_b, and C will set its frequency according to the formula:

$$f_c = 1.44/[(R_a + 2R_b)\ C\ ]$$

(with R in Ohms, C in Farads)

For example, with $R_a = R_b = 33k$ ohms, and C = .015 microfarads, the master clock frequency will be about 1000 Hz. The values shown in Figure 4 were used for the encoder and motor speed shown, they will have to be changed for a different encoder or different motor speeds. Reference [1] gives complete information on the NE555. The divider circuit uses two JK flip flops in the "toggle" mode, connected as shown in Figure 4. The 1-of-4 selector allows one of the four different speeds to be selected by the two input bits. It will select the master clock frequency (f_c), $f_c/2$, $f_c/4$, or zero. I C's 2, 3, and 5 are all common TTL chips, described fully in References [1] and [2].

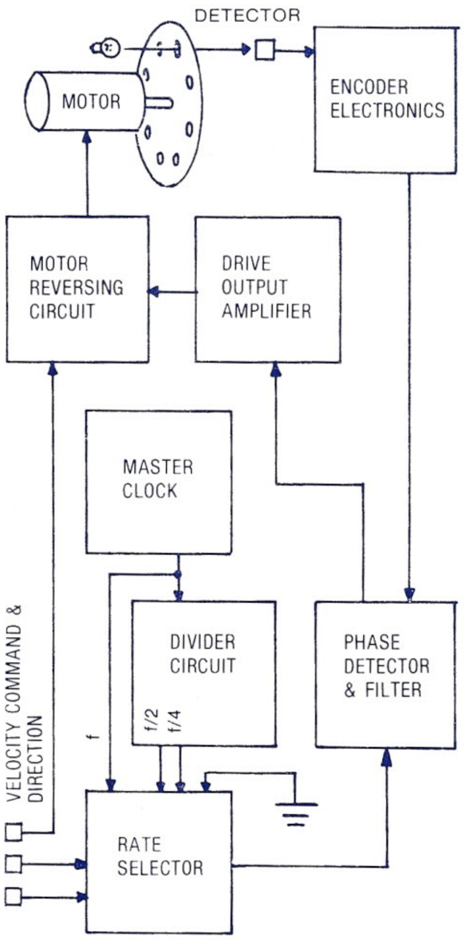

Figure 5. Block diagram of the PLS system: A three-bit velocity command is issued to the PLS system to choose one of four speeds in either the forward or reverse direction.

The MC4044, described in Reference [3], is used as the phase-frequency detector. An active low-pass filter is constructed using the amplifier in the MC4044 package and the resistors R_1, R_2, and capacitor C. Their values will set the break frequency of the filter and were found experimentally. They should be selected so your motor runs smoothly and doesn't oscillate. The values shown will probably be fine for many small motors. When experimenting with different values, try to keep the ratio R_1/R_2 about equal to 10. If the ratio is less than 10, the filter begins to pass high frequencies. The value of C will have an effect on the "damping" of the system. In general a large C will tend to make the system sluggish and take too long to lock, while too small a C will make it oscillate. The voltage at the output of the filter should vary between 0.75 and 2.25 volts. For a 12 volt motor, a gain in the driver amplifier of about 5 should allow for a good range in motor speed. The voltage gain of the driver amplifier can be changed by varying resistors R_c and R_e. The gain is approximately equal to $-R_c/R_e$. A control theorist would have a good theoretical understanding of the influences of changing the filter's characteristics and varying the loop gain, but all that's really needed is a little experimenting to find the values which best suit your motor/encoder combination. Readers interested in learning more about the control theory involved should pursue References [4], [5], and [6].

III. Operating Limitations

The circuit described here was designed to demonstrate the PLS principles and to give you a working circuit with a low chip count that can be easily added to your robot. It is capable of operating at only three different speeds in each direction, and the speeds must be 0, 1, 2, or 4 times the minimum speed. You can increase the range of possible speeds by building a more elaborate programmable reference clock circuit. Without going into complete details, several options are available that you might want to try.* If you just want particular speeds other than the ones

* *Editor's note: Readers interested in more details on these options should write us at ROBOTICS AGE. We will try to respond to your inquiries in future issues.*

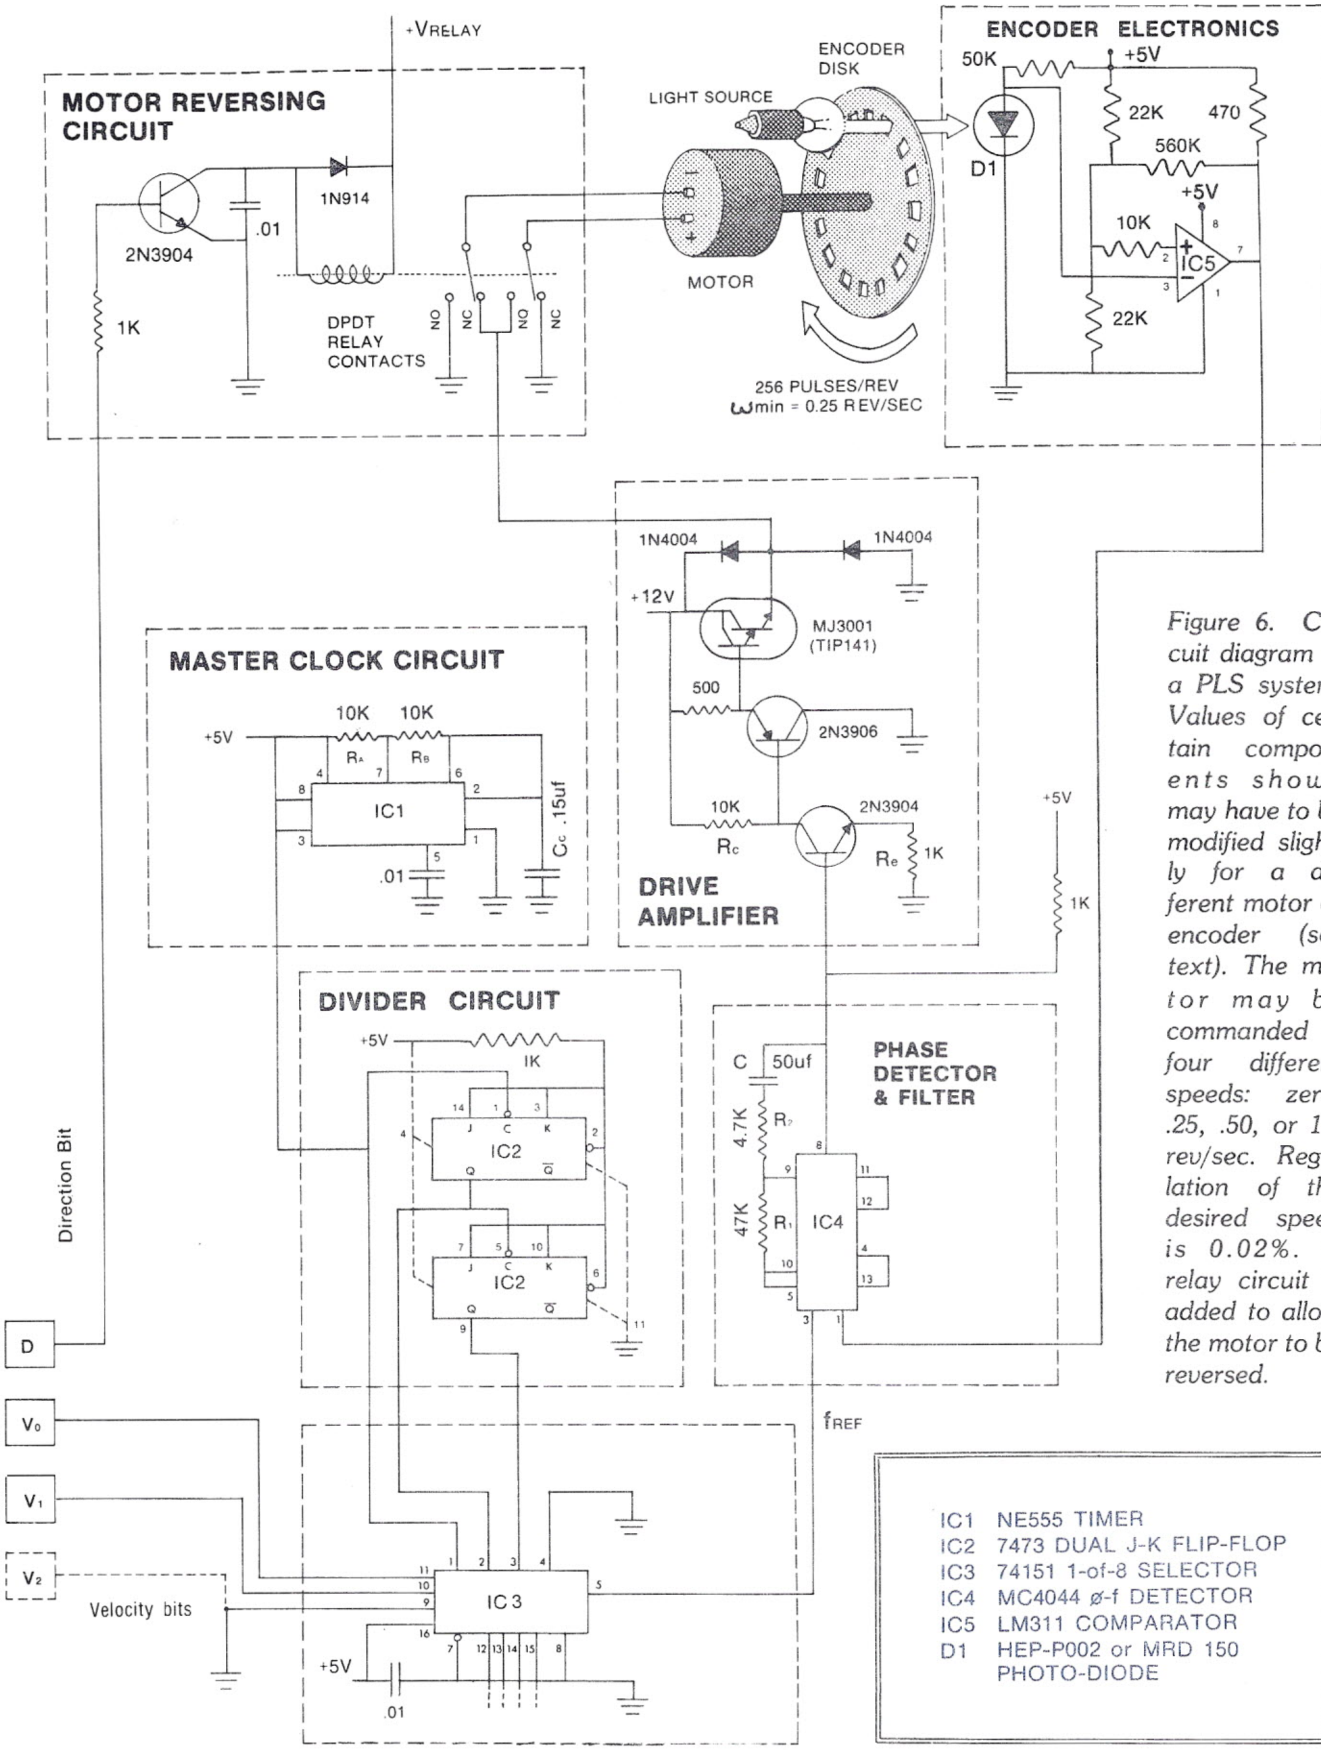

IC1 NE555 TIMER
IC2 7473 DUAL J-K FLIP-FLOP
IC3 74151 1-of-8 SELECTOR
IC4 MC4044 ø-f DETECTOR
IC5 LM311 COMPARATOR
D1 HEP-P002 or MRD 150 PHOTO-DIODE

Figure 6. Circuit diagram of a PLS system: Values of certain components shown may have to be modified slightly for a different motor or encoder (see text). The motor may be commanded to four different speeds: zero, .25, .50, or 1.0 rev/sec. Regulation of the desired speed is 0.02%. A relay circuit is added to allow the motor to be reversed.

available here, the outputs of different clocks could be selected by the input commands. The 74151 shown in the circuit is capable of selecting one of eight different inputs; for simplicity, only four inputs were used in this circuit. A circuit with the same chip count could provide eight possible speeds by using the outputs of a 74160 binary

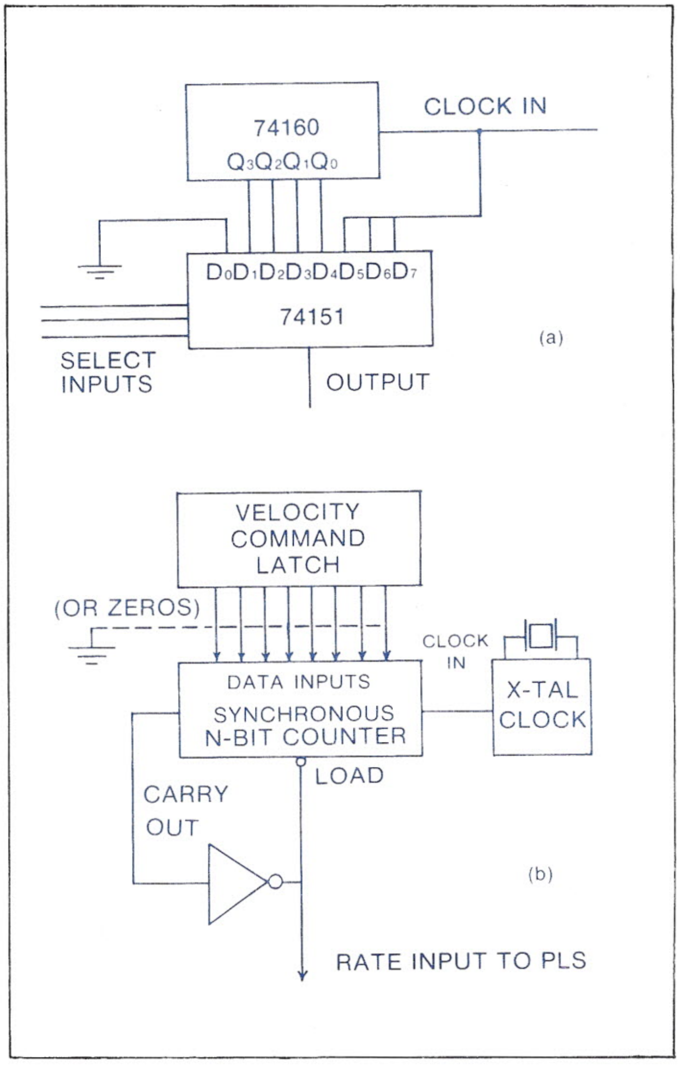

Figure 7. Options for multiple-speed selection: (a): Outputs of binary counter are used as rate outputs. (b): Carry output of counter is used as rate output and counter reload, forming a divide-by-N circuit. Rate output is clock frequency divided by (2^r-V), where r is the number of bits in the counter and V the number loaded into the counter from the velocity command latch.

counter to divide down the clock, as shown in Figure 7(a).

You can obtain a continuous range of possible speeds by using a variable frequency oscillator for the reference signal. Options here would be to use a pot to control the frequency of the 555 (Reference [1]), or to use a Voltage Controlled Oscillator (VCO) chip (Intersil 8038, or others). A disadvantage of going to an analog circuit for the clock, however, is that you lose the accuracy of the PLS; if you don't know the accuracy of the reference signal, the fact that the PLS can lock to it within 0.02% doesn't mean much. This may be fine if your robot isn't doing its own navigation. Also, if your PLS is to be interfaced to a microcomputer, you have to use a D/A converter to set the control voltage.

An accurate programmable digital clock with a practically continuous output frequency range can be built by using a crystal-controlled oscillator as the input to an N-bit binary up-counter (Figure 7(b)). Use the counter's carry output to produce the output signal and to reload the counter from the speed command latch register. This forms a "Divide-by-N" circuit with a minimum frequency of $f_o/2^r$, where f_o is the crystal clock rate and r is the number of bits in the counter. If needed, gate some zeros into the most significant bits to limit the maximum rate and define the useful range of the clock.

When given a new speed command, the reference rate suddenly becomes different from the feedback signal, and the PLS compensates by changing the motor voltage appropriately. The problem is that the motor cannot change speeds instantaneously, and during that time the velocity is not accurately controlled. The accelleration/deaccelleration curve of the system is influenced by the characteristics of the low-pass filter as determined by R_1, R_2, and C. Finding the lowest value of C that will work without oscillating gives a "critically damped" system that will lock to the reference the quickest. This is especially important when the robot is navigating; the velocity errors result in position uncertainty. Angular error in a turn causes position error proportional to the distance traveled! A solution is to make speed changes in small steps, so that the PLS only has to make small adjustments. Depending on the resolution of the speed command, a microcomputer-controlled PLS can have a programmed accelleration curve for each motor that minimizes the errors.

With a proper heatsink, the motor output transistor shown in the circuit is good for small motors requiring only a few amps at 12V. If you try to drive a more powerful motor, this drive amplifier may overload and die, since the output transistor operates in the linear region and may

dissipate too much power. The best way to control large motors is with a Pulse Width Modulated (PWM) switching power circuit, which turns the output transistor completely off and on at a controlled rate, supplying the desired average power to the motor and dissipating the minimum possible power in the transistor.*

---

* *Editor's Note: Stick with ROBOTICS AGE for details of high-power switchers and all your other robotic circuits.*

REFERENCES:

[1] Linear Data Book, National Semiconductor, 1976

[2] The TTL Data Book, Second Edition, Texas Instruments Inc., 1976

[3] Motorola Data Sheet, MC4344/MC4044 Phase Frequency Detector, MTTL Complex Functions, Issue A, Motorola

[4] Dorf, Modern Control Systems, Addison Wesley, 1974

[5] More, A. W., "Phase Locked Loops for Motor Speed Control", IEEE Spectrum, April 1973

[6] Tal, J.,"Speed Control by Phase Locked Servo Systems New Possibilities and Limitations", IEEE Transactions on Industrial Electronics and Control Instrumentation, February 1977

The author, John Craig, has experimented with robots since high school. Formerly with the Robotics Project at Rensselaer Polytechnic Institute, where he received his Master's degree, he is now a member of the Jet Propulsion Laboratory Robotics Research Program, which is developing advanced robot systems for NASA. John is one of those fortunate people whose vocation is also his avocation, and his microcomputer-controlled home robot employs many cost-saving techniques such as the PLS.

Industrial Robotics '79

Each year robotics researchers, manufacturers, users and potential users from around the world gather at the International Symposium on Industrial Robots (ISIR) to see the newest products and the latest research results in this rapidly changing field. Each year more advanced sensory and control techniques are brought out of the laboratory and into practical application to increase the capabilities and productivity of commercially available systems. Fifty-one technical papers were presented and published in the Proceedings, describing a variety of innovations, some of which will be briefly discussed here.

The presentations may be (loosely) grouped into several categories:

- robot vision systems and tactile sensing,
- improvements in mechanical design,
- computerized servo control techniques,
- programming languages for robot control,
- software tools for robot design and evaluation.

Several of the presentations discussed the application of microcomputers to sensory and control problems, emphasizing how the current semiconductor boom is resulting in increased power and flexibilty in robot systems at substantially lower cost compared to earlier methods.

Figure 1. This photo shows a robot manipulator by ASEA of Sweden. This arm, used in many European installations, uses internal push-pull rods to transmit power from its base, thereby reducing arm inertia. The robot is shown deburring a curved casting, using complex adaptive servo control. (See text.)

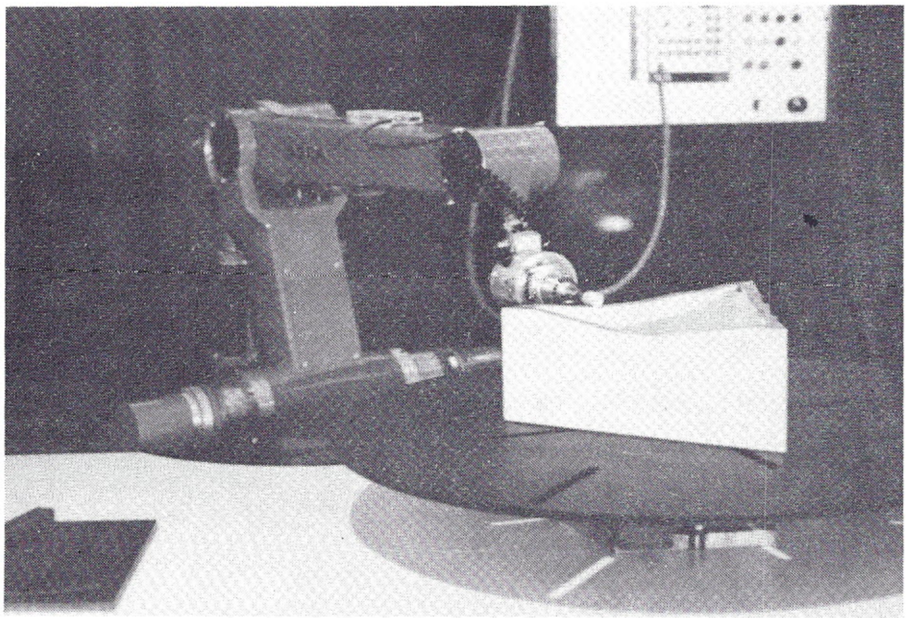

Robot Vision and Tactile Sensing

General purpose computer vision capable of recognizing objects in an unconstrained three-dimensional environment remains among the difficult challenges faced by Artificial Intelligence research. By selecting from the results of this research those techniques with high performance/cost ratios suitable for application in the industrial environment, several fast vision systems have been produced that may be economically feasible for commercial application. Researchers from most of the countries participating in the symposium described such systems. Most of the systems rely on high contrast between the part to be detected and the background — specially colored conveyor belts, light tables, etc. are used to achieve this contrast.

A vision system by SRI

International employs specially designed image processing hardware to relieve the computation burden of its integral microcomputer, resulting in rapid part location/identification times. The system consists of a General Electric solid-state camera, the image processing hardware designed by SRI, and a DEC LSI-11 microcomputer. Functioning as an independent module, the system receives commands from and communicates results to a controlling computer which can use the results to guide a manipulator, etc. It can be trained to distinguish between several kinds of parts at a known distance and perspective without special programming by placing the system in a "teach" mode and presenting an object in each of its stable resting positions. Geometric features of the object are computed and stored which can later be used to recognize the object. The system requires high contrast to allow the outlines of objects to be detected by the use of brightness threshold

Figure 2. A large hydraulic robot arm by Cincinnati Milacron, Inc., demonstrating its ability to track a moving conveyor belt and unload parts without stopping the belt. (See also Fig. 5). This arm has an integrated safety system that uses photoelectric sensors to detect intrusions in the workspace. In such a case, the arm will automatically retract to a safe position and stop.

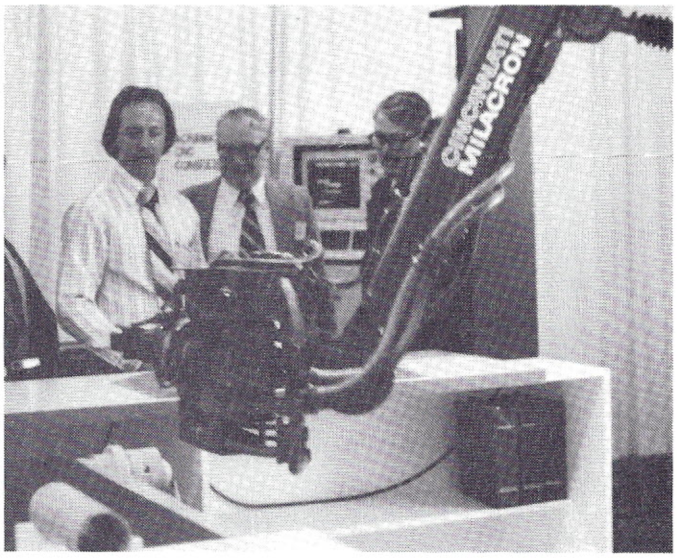

Figures 3 & 4. The PUMA arm by Unimation, Inc., which is driven by electric motors. Each joint is feedback controlled by its own microcomputer, which receives its movement commands from a controlling (DEC LSI-11) microcomputer. Designed for light industrial applications, the arm may be programmed for a variety of complex assembly tasks by expressing the problem in a robot control language called VAL. (See text.) The robot is shown inserting lightbulbs into the back of an automobile dashboard. A vision system supplied by SRI International was integrated with the PUMA controller. The system was able to locate and identify letters placed at random on a light table, enabling the PUMA to spell its name.

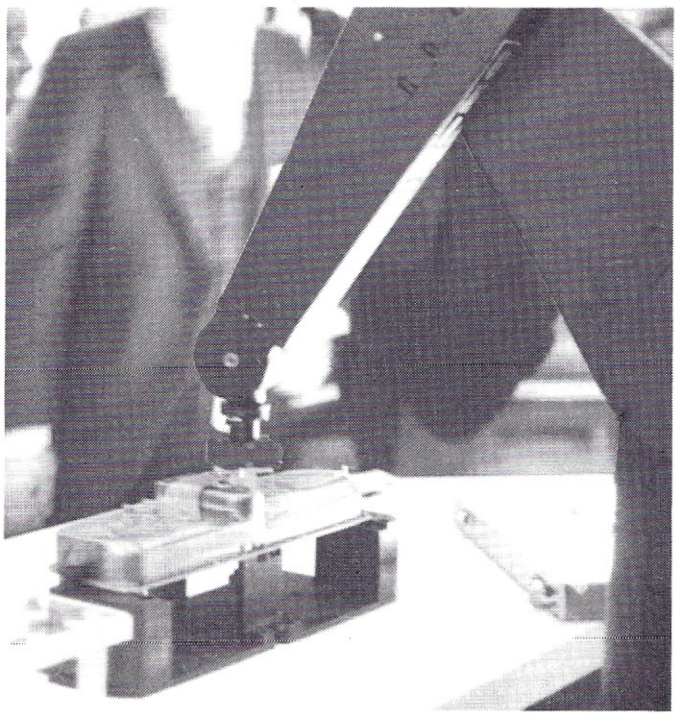

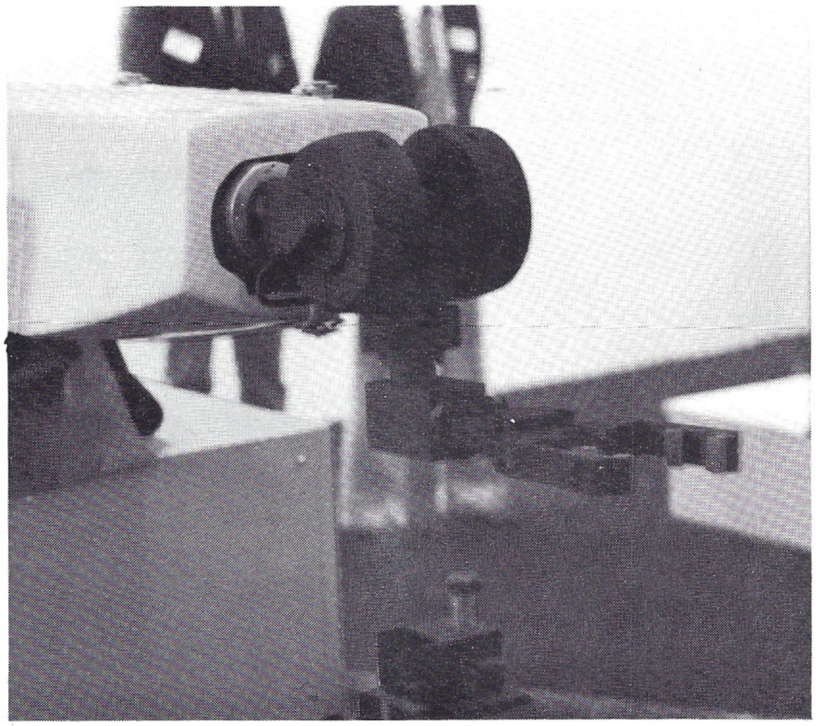

Figure 5. Many industrial robot systems require that parts must be stationary when presented for grasping by the robot. This assembly-line robot by Unimation, Inc., of Danbury, Conn., demonstrates an innovation now offered by several major industrial robot manufacturers. A rate encoder that measures the speed of the conveyor belt is included in the arm's control circuit in such a way that the belt movement is automatically added to the movements of the arm. The result is that the arm is able to behave as if all operations are being conducted in a coordinate frame at rest with respect to the belt. The need to stop the belt at the workstation is eliminated, and with it the increased cost and loss of efficiency previously due to synchronizing belt movement with workstation activity. A similar capability was demonstrated by Cincinnati Milacron (Fig. 2).

values that exclude the background. Object detection/location times of about one second were reported, with object recognition requiring longer. The system was demonstrated at the Symposium by determining the identity, location and orientation of a set of four letters placed on a light table, operating under the direction of a Unimate PUMA robot control computer. (See Figure 4).

The General Motors Research Laboratory reported a vision system that relies on carefully structured lighting to accomplish object outline detection, rather than a structured contrast with conventional lighting. A part moving on a conveyor belt interrupts a focused line of light being watched by a one-row solid-state image sensor. Again, recognition is based on 2-dimensional geometric properties of the outline. Another high-speed, high-contrast vision system by the BBC Brown Boveri Research Center in Switzerland reported cycle times of less than 200 ms. (including the time required to recognize the object as one of up to 16 stored patterns) by the use of special image processing hardware.

Most of the systems reported were limited to detecting/recognizing isolated parts. Two systems presented at the conference dealt with the problem of identifying an accessible part in a pile of known parts, the so-called "bin problem". The Robot Research Group at the University of Rhode Island discussed a system which can acquire parts from a bin by using a vacuum suction gripper. A binary thresholded image from an arm-mounted camera is used to locate candidate grasp regions larger than the sucker face, and a grasp is attempted. The flexible sucker can accomodate to the actual surface angle of the part. If the grasp attempt fails, the system will automatically retry a limited number of times on different parts. If the grasp is successful, the robot can insert the piece into the workstation at the required orientation. The total cycle time for imaging and part acquisition is about 20

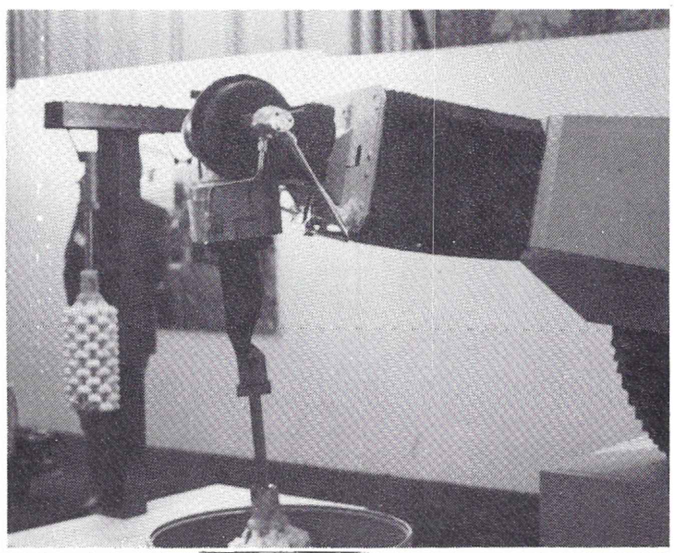

Figure 6. One of the larger Unimate manipulators is shown dipping parts to apply an industrial coating. The robot is able to acquire parts from the rack in the background and deposit dipped parts back in the rack.

Figure 7. An air-driven manipulator by PRAB, holding up a 100 lb. transmission casing. The casing is being pressed horizontally against a stress gauge. A similar gauge is mounted on the opposite side (behind the man in the foreground) and the manipulator was able to press the casing against the opposite gauge with the same amount of force.

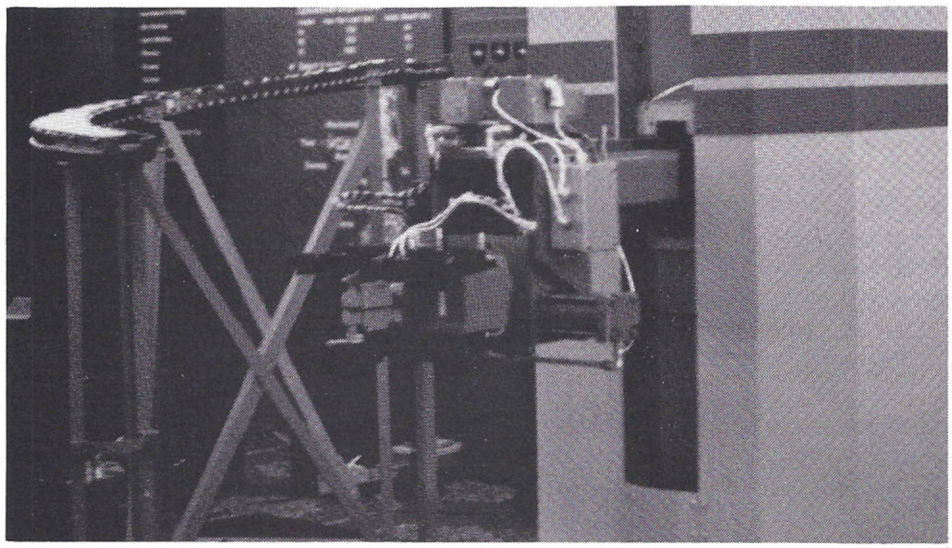

Figure 8. A "Pick and Place" robot by PRAB, Inc., of Kalamazoo, Mich. whose arm linkages form a cylindrical coordinate system. The end-effector includes dual grippers, and is demonstrated removing rods from the curved conveyor on the left. This pneumatic arm is capable of moving between manually-set fixed stops.

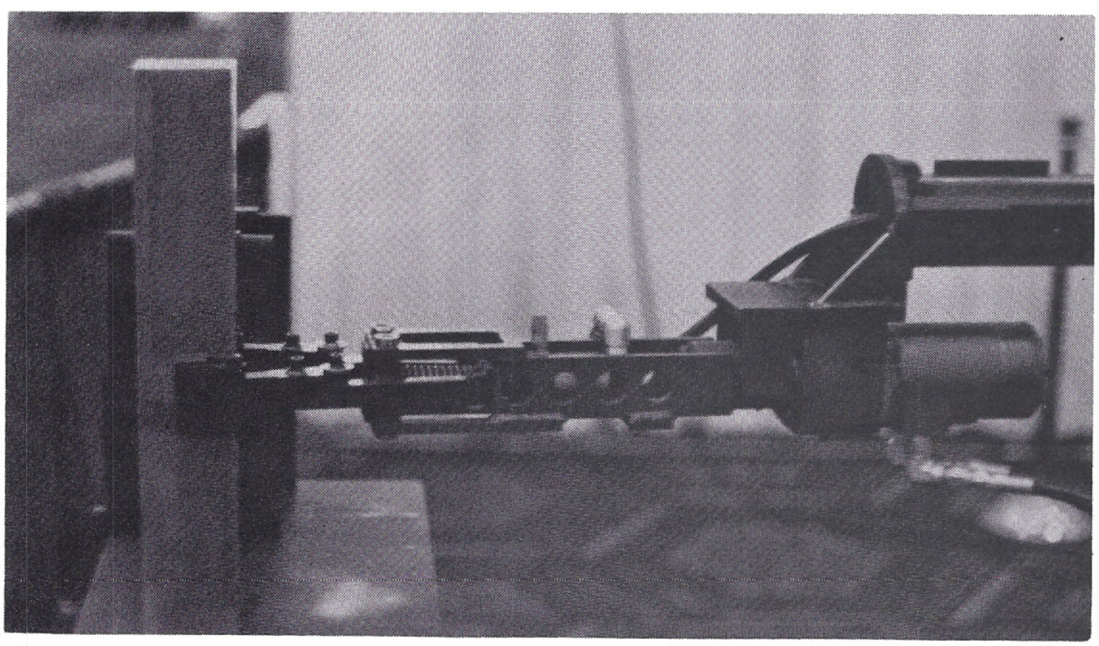

Figure 9. A robot end-effector by PRAB, demonstrating its ability to accurately position a 100 lb. weight.

seconds. The incidence of failed grasp attempts varys with the type of workpiece, with success rates from 60 to 100 percent reported for test pieces.

A system for recognizing the topmost of a pile of overlapped parts was described by researchers from the EPFL, Lausanne, Switzerland. This system required that the parts be flat and have high contrast with the background. Although this is not the actual "bin problem" overlapped parts often occur in conveyor belt feed situations satisfying these criteria. Selection of the topmost part is accomplished by comparing edge contours in the processed image with the stored outline of the object. The image contours with the highest correlation with the stored outline indicate the object on top of the pile. (A similar system developed by the GM Research Laboratory is reported elsewhere in this issue.) Processing times vary with the length of the object's perimeter, but require several seconds for test objects (using a DEC PDP-11/40 minicomputer).

Robots using tactile (touch) sensors of various kinds were reported by several research centers. One system, reported by researchers from the University of Stuttgart, West Germany, combined tactile sensing with adaptive control techniques to allow a robot to grind away welding beads on automobile bodies by following the actual path of the bead. The tactile sensors provide geometric feedback about the size and path of the bead as well as the surface smoothness, allowing the controlling computer to calculate the shape of the "ideal smooth surface" as well as the grinder path. Work at the LAAS (France) was described, in which an "artificial skin" is used to locate and recognize parts based on their pressure distribution over the sensitive surface.

Improvements in Mechanical Design

The "Remote Center Compliance" (RCC) method developed at the MIT Draper Laboratory has received considerable attention for its spectacular performance on tasks requiring a robot to insert a workpiece into a hole of close tolerance. Basically, compliance in a system allows it to appropriately respond to contact forces between the workpiece and the fixture in such a way that positioning errors are automatically corrected. RCC is a method of passive compliance affected by flexible structure in the robot's end-effector, so that errors are accomodated by the structure itself rather than by corrective action of the control system (as in active compliance). Examples of the applications of RCC both in workpiece assembly and in tooling interfaces were reported. The inventors emphasized that the method can be employed in configurations with a non-vertical insertion direction by appropriately counteracting the deflection of the compliant

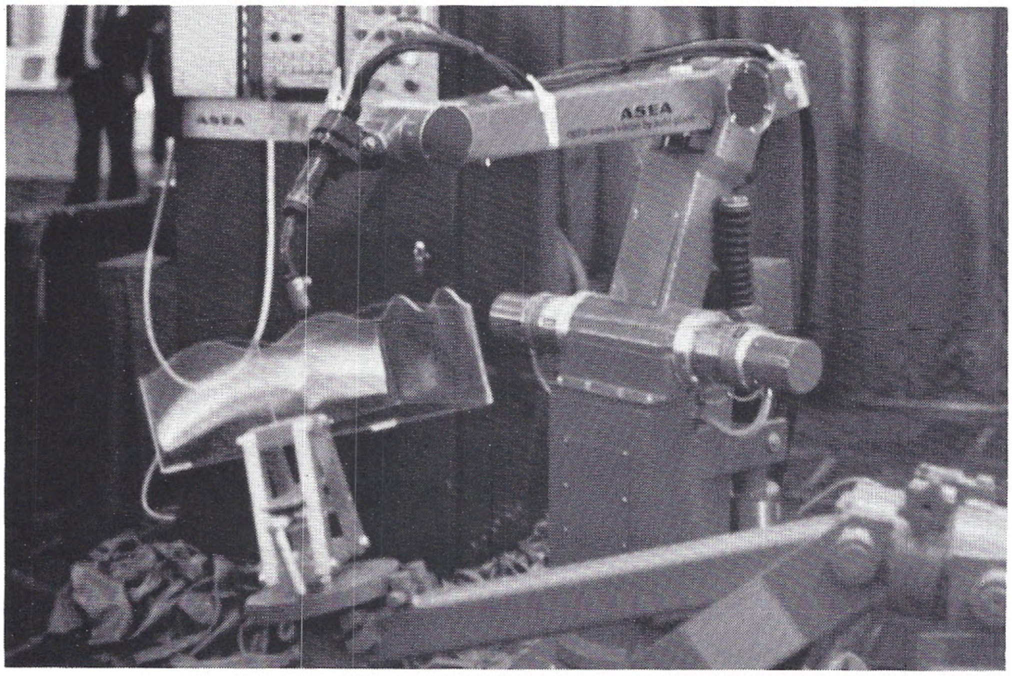

Figure 10. Another arm by ASEA demonstrated automatic arc welding. A moveable part holder is shown in the foreground.

structure caused by gravity. Passive compliance achieved by the use of elastomer laminates was also discussed.

Several other improved end-effector designs were reported. Self-aligning grippers were described by representatives of Fujitsu Fanuc, Ltd. (Japan) and Warsaw Technical University (Poland). Dual-gripper end-effectors capable of unloading a machine tool of a finished piece and inserting raw stock with one arm movement were presented by Fujitsu Fanuc and Cincinnati Milacron (USA). Milacron also introduced a completely new design for a flexible robot wrist, using concentric driveshafts driving conical gears, that can position a payload along any tangent line of a partial sphere centered on the end-effector. Driving the concentric shafts by remote actuators results in minimal workspace intrusion, and the design allows continuous reversible rotation of the payload roll axis.

Computerized Servo Control Techniques

Four papers were presented describing techniques for the control of robot manipulators by microcomputer networks. Because of the potential cost reductions possible with microcomputer-based control systems, this field has been receiving considerable attention, and the research reported at the symposium employed several different hardware organizations. Robot systems by the National Bureau of Standards (USA) and IEE of Milan (Italy) discussed network architectures in which individual microcomputer units (MCU's) are assigned either a sensory processing, control, or arithmetic function. In the NBS system these functions are organized to form sensory and control hierarchies, functioning in more abstract terms at the top levels and in actual sensor and actuator signals at the lowest functional level. Each level of the control hierarchy may use feedback from the corresponding level of the sensory hierarchy. All MPU's share a common communication bus, and may communicate by a shared memory module on the bus. In the IEE system, a communication bus is used for passing messages between MPU's and a central control processor.

The other two presentations described control algorithms for use in systems where each joint has a dedicated MPU controller. A system developed at North Carolina State University (USA) employs a simplified velocity control equation, enabling the use of a cheap, relatively slow MPU (Motorola M6800) for joint control, but still with "reasonably good" performance. A method of obtaining coordinated high-speed movements of robot manipulators was presented by researchers from Purdue University (USA). At periodic intervals during the motion between "set points" along a planned trajectory from one

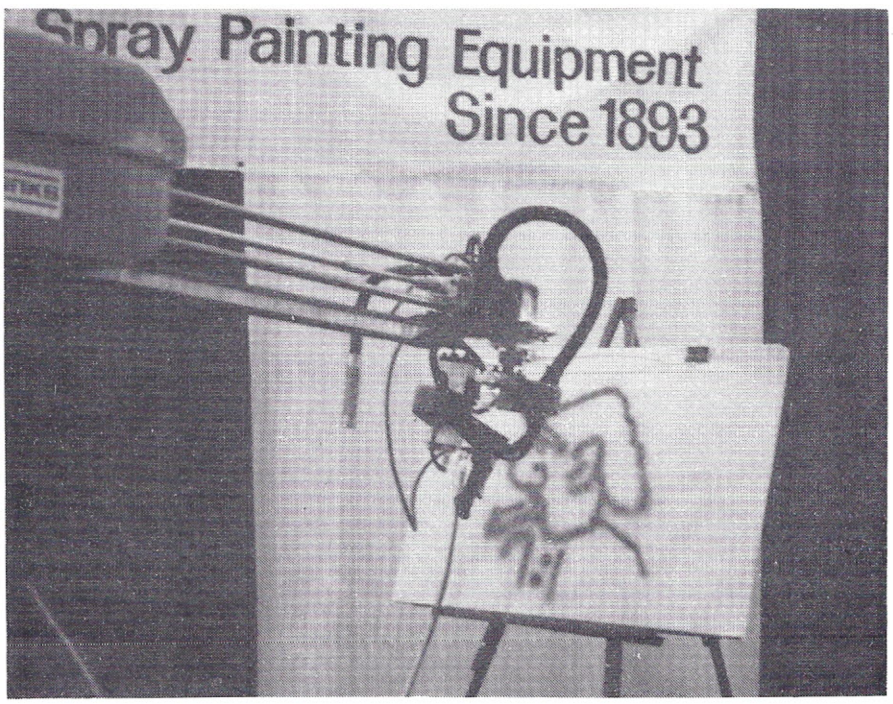

Figure 11. The Binks spray-painting robot is taught its job by leading it through the motions manually with power off to actuators. The system makes a recording of the movements and is later able to repeat the original motion automatically.

point in the workspace to another, the joint having moved the smallest fraction of its total required movement is selected as the "controlling joint" of the entire manipulator. Coordination is achieved by making the other joints the "slaves" of the lagging joint, and introducing appropriate corrective terms into the control equations of the slave joints. This approach eliminates the need for numerically modelling the trajectory in advance to compute torque curves for each joint. The method assures that even if the manipulator is driven at top speed, the rates of the faster joints will be coordinated to the slower so that the arm will pass through the set points that define the trajectory.

Papers describing several systems employing "adaptive" control techniques were presented. The term "adaptive" refers to the ability of the robot to respond appropriately to changes in the environment and still be able to carry out its task. The University of Stuttgart system (discussed above under tactile sensing) is able to follow the path of a welding bead on an automobile body and "feel" the surface for optimal grinding. Similarly, the ASEA robot, actually demonstrated at the Symposium (see Figure 1) uses an inductive sensor built into the tooling interface to follow the edge of a workpiece without either a detailed pre-planned trajectory or an exact model of the piece. The system is able to choose the optimum deburring speed, adapting to the size of the burr and slowing only when necessary for large burrs. Compared with pre-programmed systems that require complete workpiece models and which must limit deburring speed to that required for the largest burrs, this technique reduces both setup time and average cycle time. Even if the initial programming is coarse, the sensory feedback enables the system to automatically adapt its program to the actual shape of the workpiece, allowing the robot to work on parts of similar shape without reprogramming. ASEA plans to market the system this summer (1979).

An experimental system at SRI uses a small solid-state TV camera mounted on the manipulator end-effector to achieve "visual servoing". In most cases, a structured beam of projected light is used to enable rapid identification of the target feature on the workpiece. Error terms computed by the vision system are fed into the arm control loop at a rate determined by the image processing cycle time (about 150-500 msec. using a DEC PDP-11/40). Several experiments were reported in which a Unimate arm controlled by visual servoing was able to track an object on a conveyor belt of unknown velocity, perform a simulated spot weld on a part on the moving belt, and to follow a curved path in three dimensions at a constant velocity, simulating such tasks as gluing, seam-following for welding, etc.

Programming Languages for Robot Control

Computer control of robots is rapidly replacing the various hard-wired or fixed-sequence methods previously used, resulting in greater flexibility and power in industrial applications. Ease of programming and the ability to integrate sensory input into the control functions are only two of the many features made possible by using computer control. The need to program complex robot functions into the controller has led to the development of robot programming languages that allow such control functions to be expressed much more easily than by coding the problem using the native code of the machine or even in general-purpose "high-level" programming languages. Three robot programming languages were discussed at the symposium.

Rather than adding to the growing number of such systems already in use, one of the presentations described improvements to an existing robot programming language, RTL/2, made possible by creating an interactive program development enviornment that can significantly reduce the setup time for new tasks. The system, developed jointly by SRI International and Philips Research Laboratories (England) uses a subset of the RTL/2 language. Features provided by the new system include single-step operation, execution tracing, breakpoints, and data value displays. Running programs can be interrupted, either by a run-time error or manually by the programmer/trainer, and data values and/or program statements can be changed. If no active statement is changed, the system can continue the program execution from the point of the interrup without having to restart, a feature that can avoid having to perform costly setups every time a small change is made in the robot's program. The language includes many of the features of modern "structured" programming, proven effective in reducing development time and simplifying program maintenance. Other "non-structured" languages were discussed by representatives of Cincinnati Milacron, Inc. and IEE of Milan, Italy.

The utility of such robot programming languages was demonstrated quite well during the demonstration of the PUMA manipulator by Unimate. (See Figure 4.) Part of the demonstration of the system consisted of the PUMA spelling its name using letters randomly placed on a light table. The backlighted table was used to provide the silhouettes required by a SRI vision system (see above),

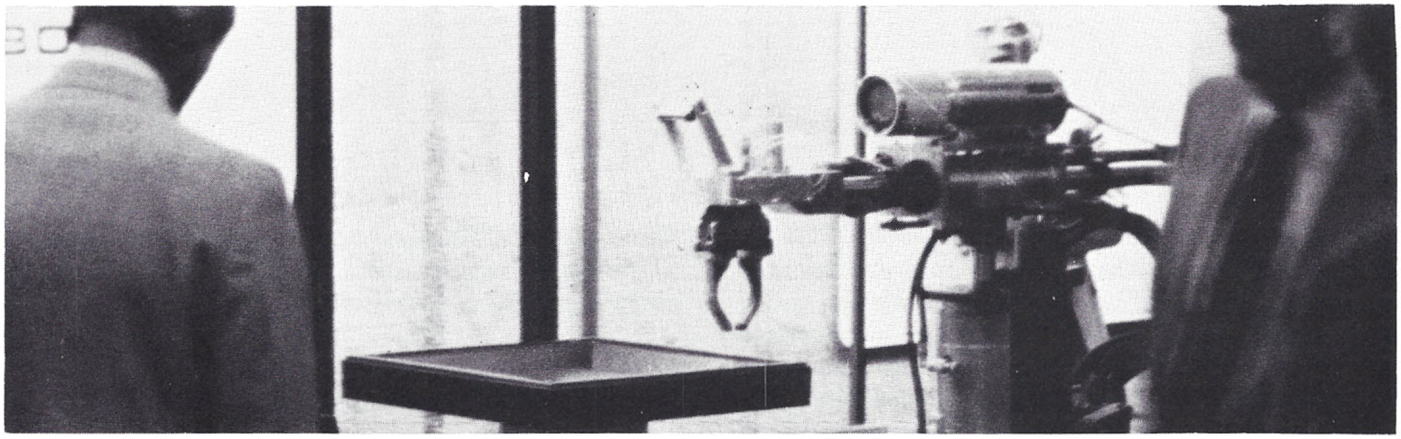

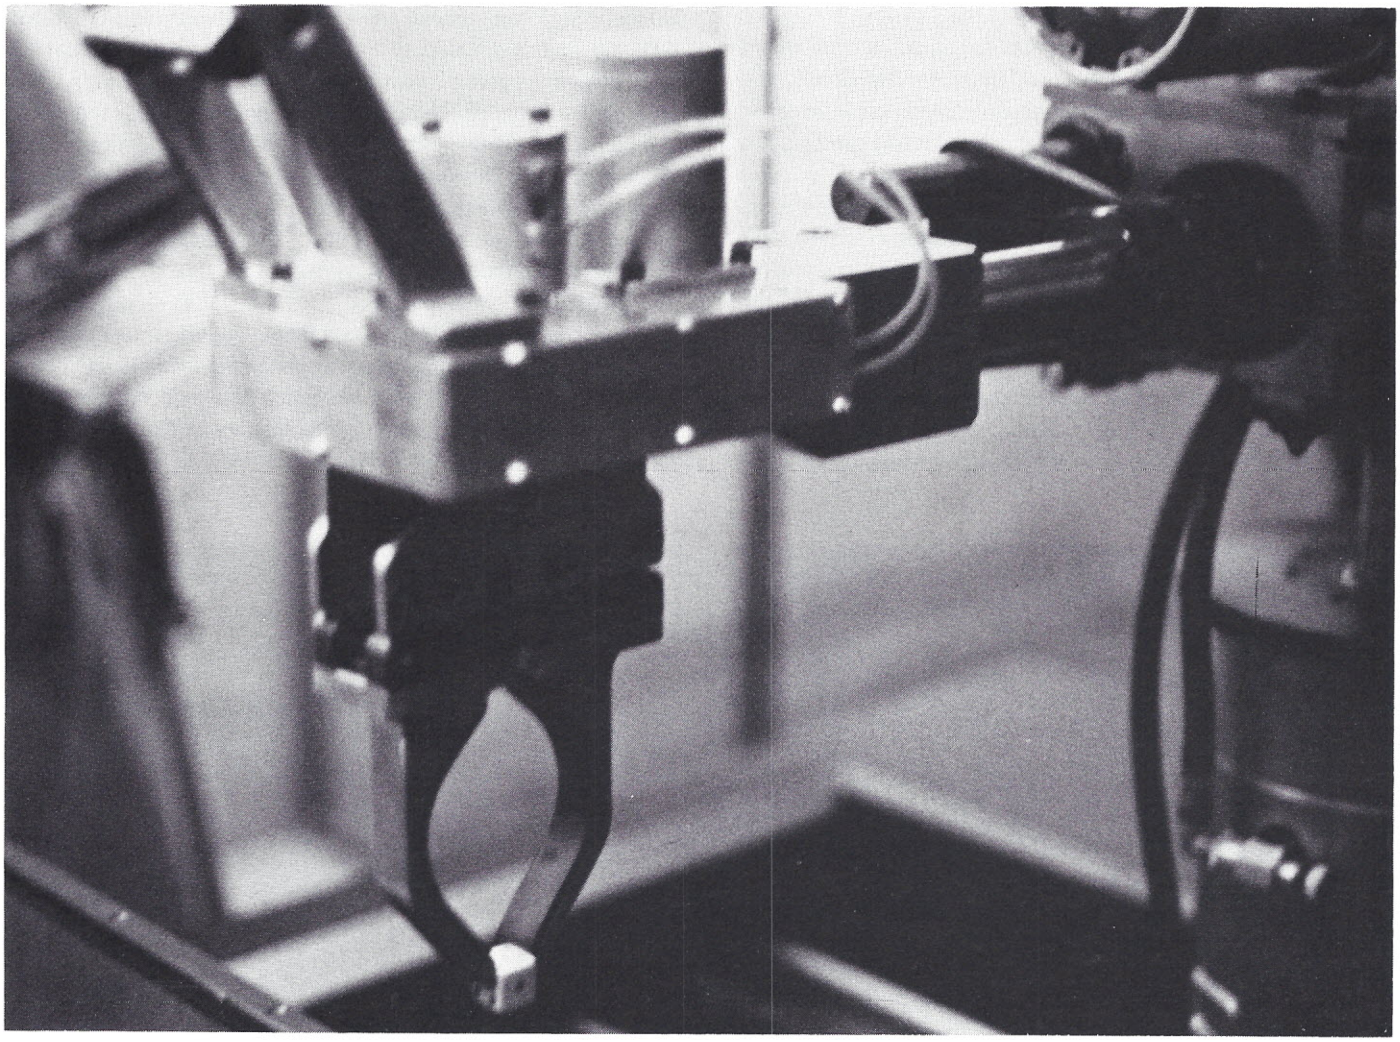

Figures 12 & 13. This pneumatic robot manipulator by Autoplace, Inc. has an integrated vision system, using cameras mounted both in the canopy above the workstation and on the arm itself. The arm camera views the work area through a mirror attached at the wrist. In previous demonstrations this robot impressed spectators by sorting a deck of cards into suits. This year it entertained them by playing dice. The vision system was able to locate the die, determine the number of the top face, and pick it up. Recognizing the pattern on the face of the die involves physically rotating the arm camera so that the pattern aligns with the stored templates, rather than performing this function by computer software. Since the Autoplace arm moves between fixed stops, an X-Y table positioned by lead screws is used to position the arm over the die once it is located by the top-view camera.

which was locating particular letters for the PUMA's control computer. The control program for the demonstration was written in Unimate's robot programming language, VAL. The program was designed to communicate with the vision system to receive the coordinates of the various letters and then line them up to spell the name. Upon completion, the Unimate representative at the demonstration was to remove the letters from the table, because the program would automatically repeat the operation using a second set of letters also on the table. At one point in the demo, the employee was engaged in conversation with an interested bystander (who told us this story) and failed to remove the last spelling before the robot repeated the cycle. The PUMA attempted to place the next spelling of its name in the predefined location, resulting in a minor collision with the (appx. 3/4" thick) letters already there. The spectator, familiar with the VAL language, suggested that the program should be revised to remove the letters itself, after completing the spelling, to make room for the next cycle. After a brief period of interactive program revision, the spelling procedure was revised to individually remove the letters "P-U-M-A" and stack them vertically off to one side of the table. No manual guidance or physical changes to the system were required, as all the new actions were easily described in software.

Software Tools for Robot Design and Evaluation

The remaining category addressed by papers in the proceedings pertained to methods of simulating robot designs and workstation configurations as an evaluation tool. Related systems were described by researchers from the University of Nottingham (England) and the MIT Draper Labs. In the Nottingham system, a designer is able to interactively design the layout of the workstation. Using computer graphics to create a representation of the workspace inside the computer, the designer can use a "lightpen" to trace the movements of the manipulator's "tool center point". After defining the sequence of operations, the computer simulates the execution of the sequence, producing both an animated display of the cycle and an estimated cycle time. Since the internal graphics model can be rotated, possible collisions in the workspace can be determined by viewing the motion of the model

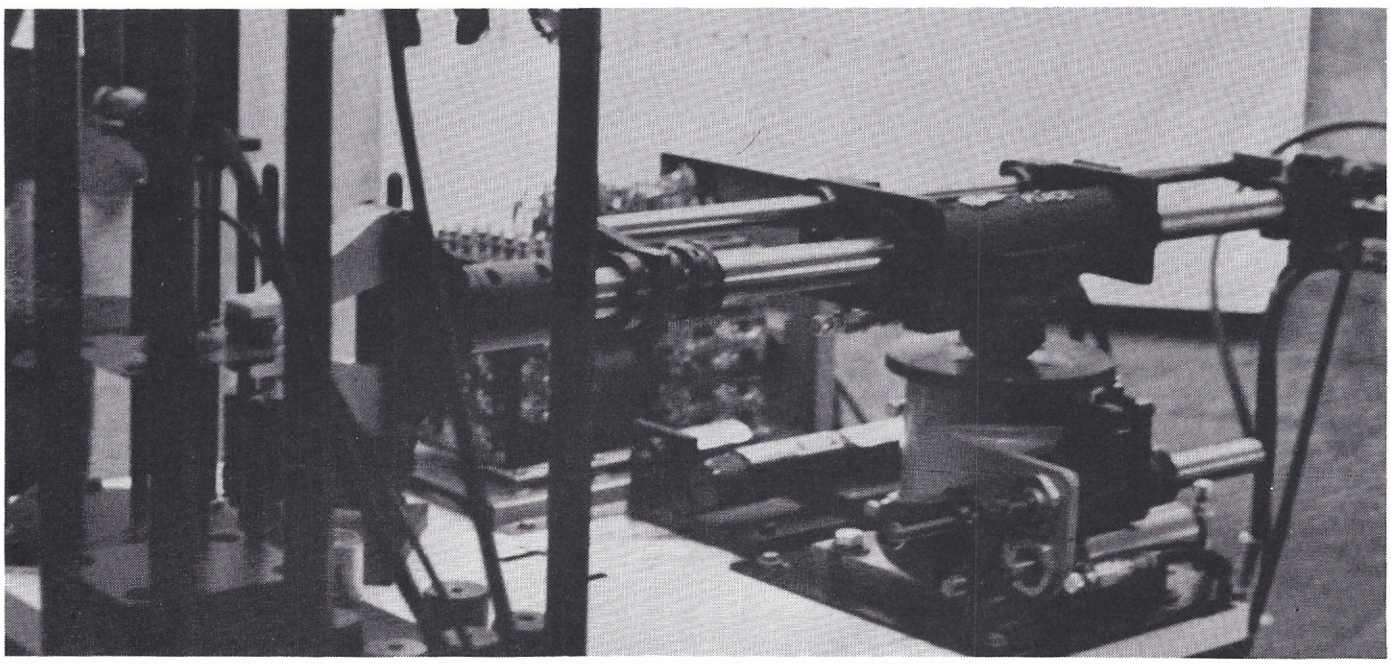

Figure 14. Robot arm by Autoplace, Inc., of Troy, Michigan. Parts removed from a feeder in back are held up for examination by a laser part verification system. The framework at left supports mirrors, configured so that when the robot holds parts at the verification station in any of several ways, an interruption of the laser beam indicates a defective part. If the part is properly formed, no interruption occurs, and the robot accepts it.

from different perspectives. Models of many commonly used manipulators and related tools are available, so that a variety of configurations may be explored with little encoding effort. In the Draper Lab system, a more detailed analysis of the phases of manipulator movement (gross, interface, and fine movements) permits more accurate estimation of cycle times, as verified by experimentation with actual robot systems. The method has been extended to analysis of multi-arm systems, and includes the effects of the failure rates of parts feeders and related equipment. Methods for developing manipulator designs based on given specifications were discussed by workers from the Institut de Microtechnique (Switzerland) and from Hitachi, Ltd. (Japan).

Summary

We have attempted to give a brief overview of the technical results reported at the Symposium. Clearly, an exhaustive categorization of every paper presented is beyond the scope of this brief article. The Proceedings are published by the Society of Manufacturing Engineers at 1 SME Drive, Dearborn, Michigan 48128.

The work presented at this year's Symposium represents the steady advancement of the state of the art in commercially available systems — the product of successful efforts by robotics researchers around the world, tempered by the practicality and cost-effective feasibility required by competitive international markets. The modern industrial robot has earned its place as a keystone of the productivity of all industrialized nations. In the words of John J. Wallace, President of the Robot Institute of America, one of the Symposium's sponsors, "It seems that we have developed an international community dedicated to the success of the industrial robot in industry, and in so doing we have avoided harmful sensationalism and built a fine record of successful performance for the industrial robot." ROBOTICS AGE applauds the innovations presented at the 9th ISIR, and looks forward to providing our readers with coverage of the ISIR's in coming years. ®

The problem of 'understanding' sensory input is one of the central themes of robotics. This is the first in a recurring series on Robot Vision, written by professional researchers in the field.

Introduction to Robot Vision

by Alan M. Thompson

Jet Propulsion Laboratory
Pasadena, California

Few people stop to consider the complexities involved in converting the physical inputs to their senses into conscious recognition of the objects or phenomena causing them. A burst of acoustical energy distributed appropriately in frequency and duration over the audible spectrum can be changed by the action of ear and mind into the perception of the word, "Hello!," said by your best friend, seemingly without any delay after it is spoken. Similarly, a two-dimensional pattern of reflected light, with appropriate variation in brightness and color, is easily recognized as your friend's face within a fraction of a second of your first glance.

The volume of processing required to distinguish each of these particular patterns from all the other possible combinations of sounds and objects that you know has so far been greater than the resources of our largest and fastest computers, programmed by our best researchers. This may be surprising to those who have the impression that all that's required to make a computer "see" is to "plug in" a TV camera. Just the volume of "raw" data contained in a standard video signal is greater than the input capacity of all but the larger minicomputers, and is enough to easily carry several *thousand* telephone conversations simultaneously (thus easily proving the old

proverb about the worth of one picture). To make the information contained in this signal accessible to a program requires suitable sampling hardware, and that is only the prerequisite for the complex process of recognition.

Humans have the ability to recognize familiar objects from almost any angle, over a broad range of distances and lighting. Most of the physiological and psychological processes that underly this ability are unknown, and, since most of recognition occurs beneath the level of consciousness, the phenomena can only be studied indirectly. With little evidence available from living systems, vision researchers in the field of artificial intelligence must develop theories that attempt to explain recognition in a way that can be duplicated in a computer. This involves deciding what knowledge is needed about the structure and appearance of objects, how it can be represented as a computer data structure, and designing procedures for using the knowledge to achieve recognition.

Numerous useful systems have been developed by restricting the recognition task to distinguishing between a limited number of objects in a carefully controlled environment. In these special cases techniques can be developed that take advantage of the simplification of the problem by being limited to the analysis of a few well chosen characteristics. These methods, some of which will be described later, have found successful applications in military and industrial systems, and offer researchers important clues to the structure of more effective general theories.

Video Image Input

To understand the process of setting a picture into a computer it is useful to refer to the "half-tone" technique used by printers for reproducing photographs. Photos are printed as a matrix of tiny dots; the size of the dot determines the brightness of the picture at that point. The photo in Figure 1 shows this process exaggerated, so that you may have to move back to recognize the subject. The size of a dot (black or white) determines the percentage of white area contained in that cell of the matrix. This percentage gives a measure of the brightness of that cell, allowing a digital representation of the image as the brightness number (referred to as the grey level) of each picture cell (referred to as a pixel) stored in a matrix that corresponds to the original picture.

Normal video, however, is an analog signal, consisting of

Figure 1. Half-tone method of reproducing photographs (exaggerated).

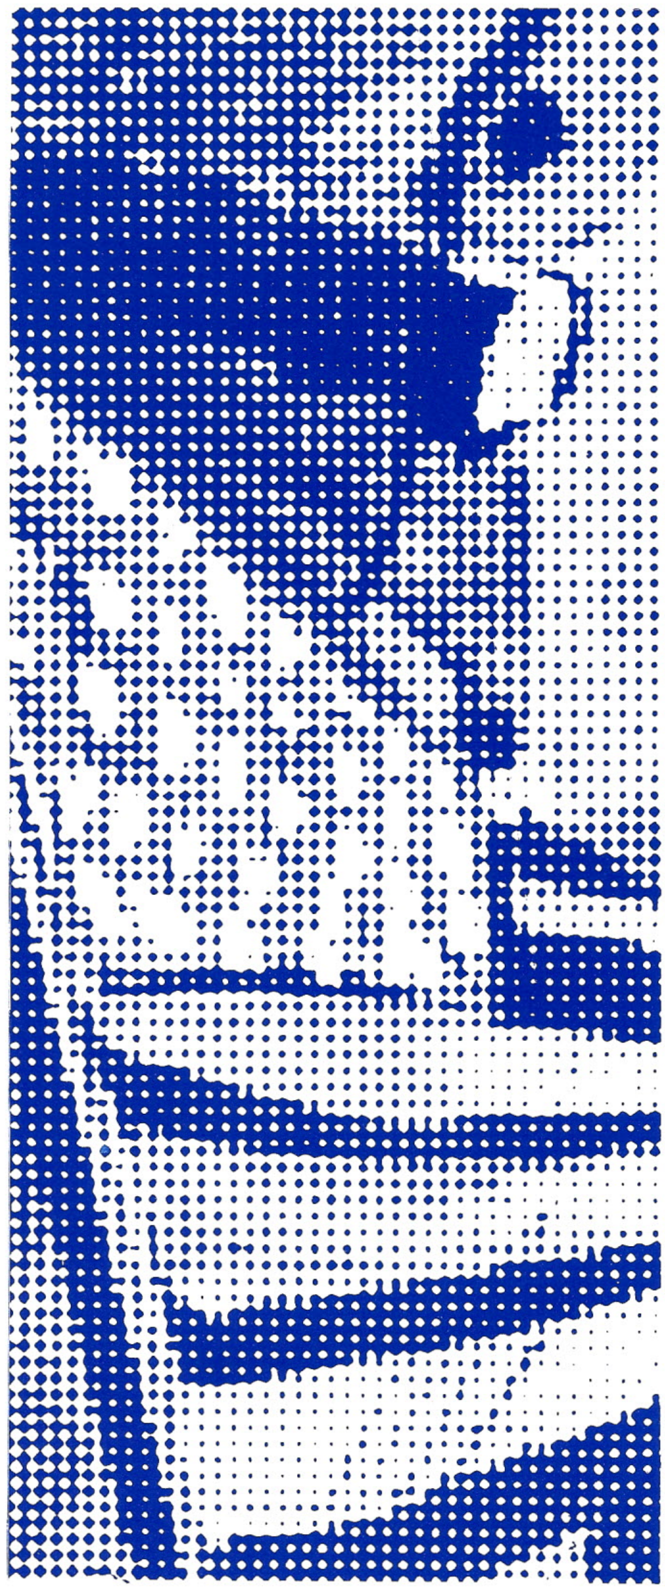

a series of horizontal "raster" scans of the image from top to bottom. Each line of the scan begins with a synchronization pulse of negative polarity, followed by a positive waveform that corresponds to the brightness of the image as a function of time as the scan moves across the image. (Color information is multiplexed into a higher band, and then decoded at the receiver.) Detection of the next horizontal synch pulse signals the end of the line and causes the scan to jump to the beginning of the next line down. (Figure 2.) When the scan reaches the bottom of the image, a vertical synchronization pulse occurs, and the scan returns to the top for the start of a new "frame". To convert this signal to a pixel matrix (for a black and white picture) a clock circuit is synchronized to a multiple of the horizontal scan frequency. The number of pulses of the "pixel clock" per scan line is the number of pixels per line in the digitized image. At each pulse, the video signal is sampled and converted to binary, and the results are deposited into the computer's memory (see Figure 3). If a sufficiently slow scanning rate is used, a normal I/O channel may be used to read the converted grey levels, but for standard broadcast video rates, the use of a dual-port video memory is essential on all but the fastest computers. To avoid distortion in the converted image, the video scan should be linear, that is, the horizontal position of the scan is proportional to the elapsed time since the horizontal synch and the vertical position is proportional to the number of horizontal lines since the vertical synch. In the new solid-state imaging arrays using CCD or CID LSI technology, each pixel position is fixed in the silicon matrix, so linearity is guaranteed. On some models, however, the horizontal pixel spacing is not the same as the vertical, resulting in a linear distortion in the converted image that can easily be accomodated by the image processing program. Also, the drive circuitry for solid-state cameras produces a suitable pixel clock, simplifying the sampling circuit.

Once the image is represented in memory as a two-dimensional array of grey-levels, the task of enhancement and analysis can begin. Image enhancement usually refers to the processing performed to improve the quality of a picture for subsequent use by humans, (as done in the unmanned space program at JPL to process spacecraft images of planets) but may also be useful as the first step of automatic image analysis. Examples of such procedures are contrast stretch and normalization, in which a transformation is computed that maps the original range of grey levels into one that more fully utilizes the available range of the grey scale. In most automatic image recognition systems, however, a transformation of the entire image is usually not necessary, since the program can use operators that are insensitive to the actual range of brightness.

Figure 2. A standard analog video signal. Each line of the raster scan is represented by a signal whose voltage is proportional to the brightness along the line.

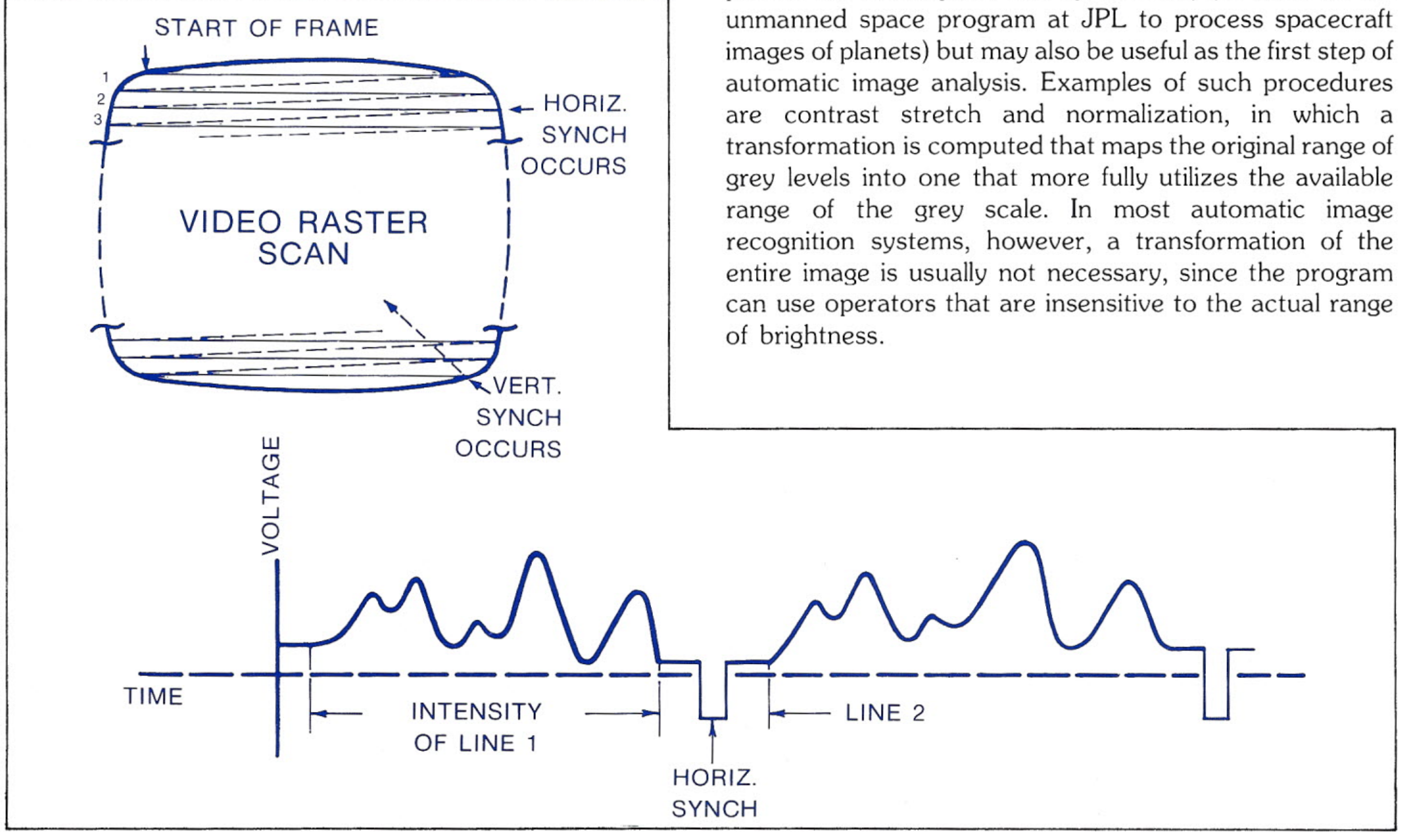

"Low-Level Processing", or Finding Primitive Features

The first step of image analysis is to locate areas in the image that correspond to whole objects or parts thereof. Ideally, there should be an easy way of isolating the image of each object from the background, (as in the case of a white egg sitting on black velvet) so that recognition techniques can be applied to just the object of interest. Unfortunately, in most environments this is not always possible. A trick used in some industrial applications is to paint the background a different color than the objects to be identified, but most research systems that operate in less constrained environments rely on methods that recognize objects without first having to determine their complete outlines. Two of the most common tools used to find objects are edge-detection and clustering. Both techniques attempt to locate the boundaries of objects or regions in the image so that their location, size, and shape may be computed as clues for recognition.

Edge detection is based on the fact that there is often a difference in brightness between object and background. Areas of high contrast can be found in the digitized image by looking through the pixel array for jumps in the grey level. If color information is available, this task is simplified because there could be a color change that might not be as noticeable based on intensity alone. Edges are located in the array by applying an operator that examines a small neighborhood of adjacent pixels and computes an "edge probability" value. If the value is above a set threshold, a notation is made in some data structure (possibly the pixel array itself) signifying the presence of an edge at the appropriate spot in the image.

The size of the neighborhood or "window" examined by the edge detector determines its sensitivity to noise. The grey level of any one pixel is subject to numerous noise sources that may result in considerable error. The accumulation of photons by the sampling element, whether vidicon tube or solid state array, is a statistical process subject to a normal error distribution based on intensity and sampling time. Signal transmission and the A/D conversion process all contribute electromagnetic noise. The presence of a noise 'spike' in the pixel array can lead to the detection of false edges. To help avoid false edges, adjacent pixels can be averaged by the edge detection function to reduce the effect of the noise. An edge in the image may be oriented at any angle in the neighborhood "window" around the point to be tested. A detector that looks for edges only horizontally or vertically loses much of the information in the picture. There is always a tradeoff between the reliability of an edge detector and its speed. Averaging over a larger window is slower and more accurate, and is capable of detecting fuzzy edges that a smaller window might miss. Detecting lines or spots requires special operators that measure the difference in brightness between a central region and its surroundings. A survey of popular edge detection techniques can be found in [1].

A typical edge detection function is the Sobel gradient operator. Given that the image is represented as an array, let the 3x3 window centered around one pixel be represented by:

A	B	C
D	E	F
G	H	I

Figure 3. Conversion from analog video to digital grey levels.

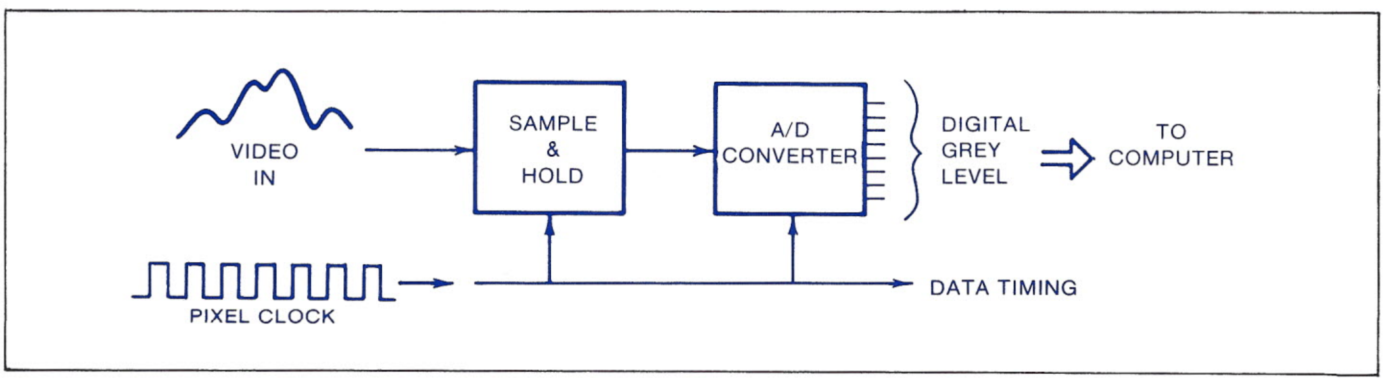

PHOTOS COURTESY OF GM RESEARCH LABS

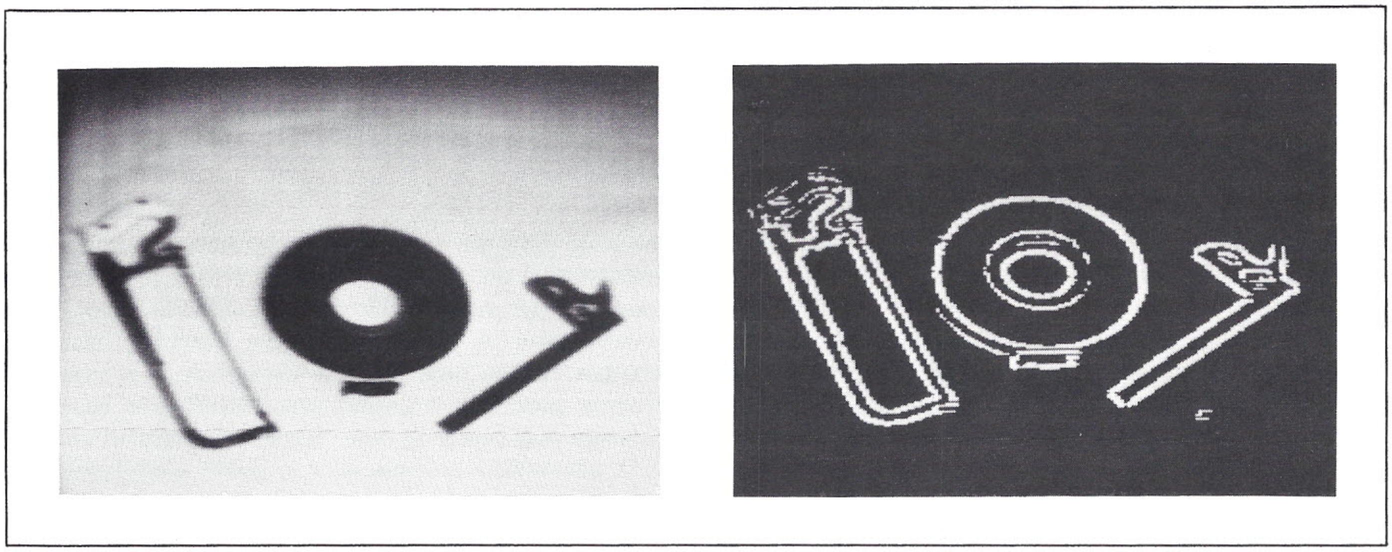

Figure 4. Edge detection. The video image (left) is processed by an edge detection operator to produce an "edge picture". (right) Processing may be done by software or, as shown here, by digital hardware operating at full video rates.

Then the 'edge value' at pixel E is computed by the formula:

$$\text{edge}(E) = \left\{ [\,(A+2B+C) - (G+2H+I)\,]^2 + [\,(A+2D+G) - (C+2F+I)\,]^2 \right\}^{1/2}$$

where, inside the square root, the letters stand for the grey levels of the corresponding pixels. Edge values above an experimentally determined threshold value are assumed to indicate the presence of the edge of some region in the picture such as an object or shadow.

This edge detection function has the advantage of being relatively fast, although the small sampling window (3x3) makes it more sensitive to noise and less sensitive to fuzzy edges. It is easily implemented in software and can be simplified further by using the approximation:

$$\sqrt{(a^2 + b^2)} \approx |a| + |b|$$

Summing the absolute values of the alternate differences avoids both the multiplications and the square root, and, because of its simplicity, can be implemented in a fast digital hardware module as part of a robot vision system. The JPL Robotics Project has such a hardware implementation, capable of producing an "edge picture" in "real time" at standard video rate, allowing the robot to track several moving objects simultaneously. [2] The basic operation of the edge detection function is shown in Figure 4. The output of the detector is displayed as the "edge picture" of the simple image shown. Note that even in this high-contrast image, there are still some gaps in the edges produced by the detector. A more complex edge detection function would produce a less fragmented outline at the expense of greater computation, but it is generally accepted that advanced object recognition programs should be capable of functioning even with incomplete, fragmented outlines of the objects to be recognized.

"Clustering", also known as "region-growing", refers to a process of collecting clusters of adjacent pixels that have similar properties. The decision whether or not to include a pixel in the region is made by a "discrimination function" which computes the desired property and compares it with a threshold value or with some value computed from the other pixels in the region. The simplest property of a pixel is its grey level, and many industrial robot vision systems locate objects merely by looking for regions whose grey level is sufficiently above (or below) that of a known background (conveyer belt, etc.) and produce a "binary" image in which each pixel is reduced to one bit, on or off according to the result of the threshold test. More complex clustering schemes use functions comparable to edge detectors that compute the property from a neighborhood, and, starting from a "seed" pixel in the interior of a region, expand the region to its boundary. Region growing has the advantage that the boundary of the region is always closed, so that it may be used to compute properties derived from its shape, such as various integral/statistical moments, etc. Noise in the image may result in inaccurate boundaries rather than in fragmented edges.

Image Recognition Methods

For many robot vision systems, the interpretation or classification of the processed outline is not necessary, since only the location and/or orientation of objects is needed (for grasping coordinates, etc.) If, however,

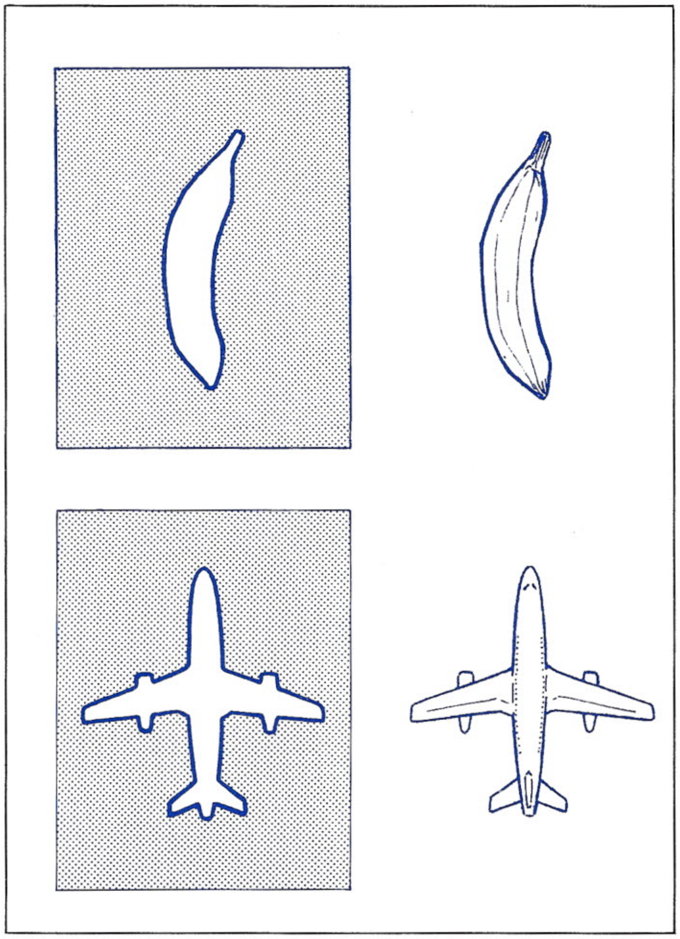

Figure 5. Template matching refers to the comparison of stored shapes with the observed image.

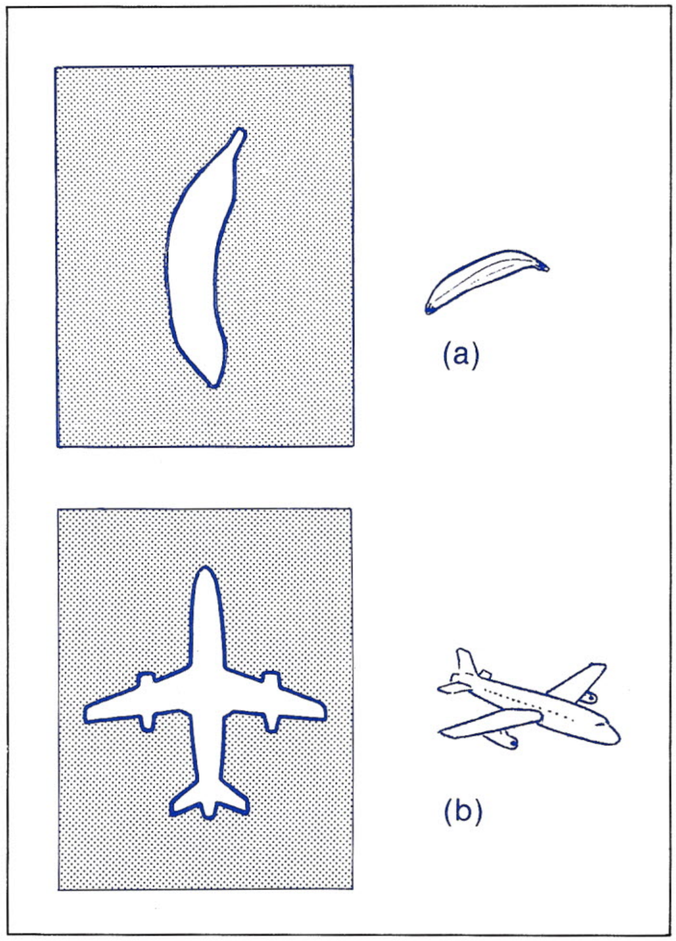

Figure 6. Problems with template matching: (a) Rotation and scaling, (b) perspective rotation.

recognition of objects is needed to achieve more "intelligent" behavior of the robot, then after objects in the field of view have been detected by one or more of the primitive operations described above, the resulting evidence is input to a cognitive process which (hopefully) determines the most likely interpretation of the image. The geometric structure of the edges, properties of the regions, and relationships between regions may be measured and compared to internal "models" of known objects. Models may take many forms, ranging from stored images, to statistical properties computed from the edges or regions, to abstract relations between the components of structures, or to various combinations of such methods. The design of a representation depends on the environment in which the system is to function, the number of objects in the database, and an analysis of their distinguishing features.

One simple recognition method that can work well in a restricted environment is "Template-Matching". This method works by comparing the image of an object (or its edge picture) with stored images (or outlines) of known objects. The stored "template" with the highest statistical correlation with the observed image is picked as the result. (Fig. 5) Clearly, this comparison is subject to many constraints. Assuming for a moment that it is possible to isolate a complete object from the background in the image, the image of that object has to be "normalized" to a form suitable for comparing to the templates. If the object can appear anywhere in the scene, then a "window" around the object must be selected that matches the template size. If the object could be rotated in the image plane to a different orientation than that of its template, the image in the window may have to be rotated to the "normalized" orientation of the template, or the template transformed into the image. Since rotation of the image would be computationally expensive, requiring that the program perform a 2-dimensional coordinate transform for each pixel in the window, most systems that require such transformation do so by transforming a vector description of the outlines. Worse still, if the viewing area is too large, the image may have distortions due to the perspective transformation of the camera that reduce its similarity to the template. Worst of all, if the object could be rotated arbitrarily in three dimensions it might have a completely different appearance, rendering any single stored outline completely useless without complex, time-consuming 3-D surface rotational transforms. (See Figure 6.)

In circumstances where objects need not be recognized

from an arbitrary perspective and in which the field of view can be suitably restricted, template-matching and related methods can perform quite well. Since the template may tell the program where in the image to look for edges, it is possible for the program to expend more effort in looking for edges where they are expected. Alternatively, if edge fragments fall on or near the outline defined by a template, a reliable match may be obtained even in cases where the edge data are imperfect due to noise, poor contrast, or the limited ability of the edge detection operator. Templates can also be used to match outlines of different sizes by

PHOTOS COURTESY OF GM RESEARCH LABS

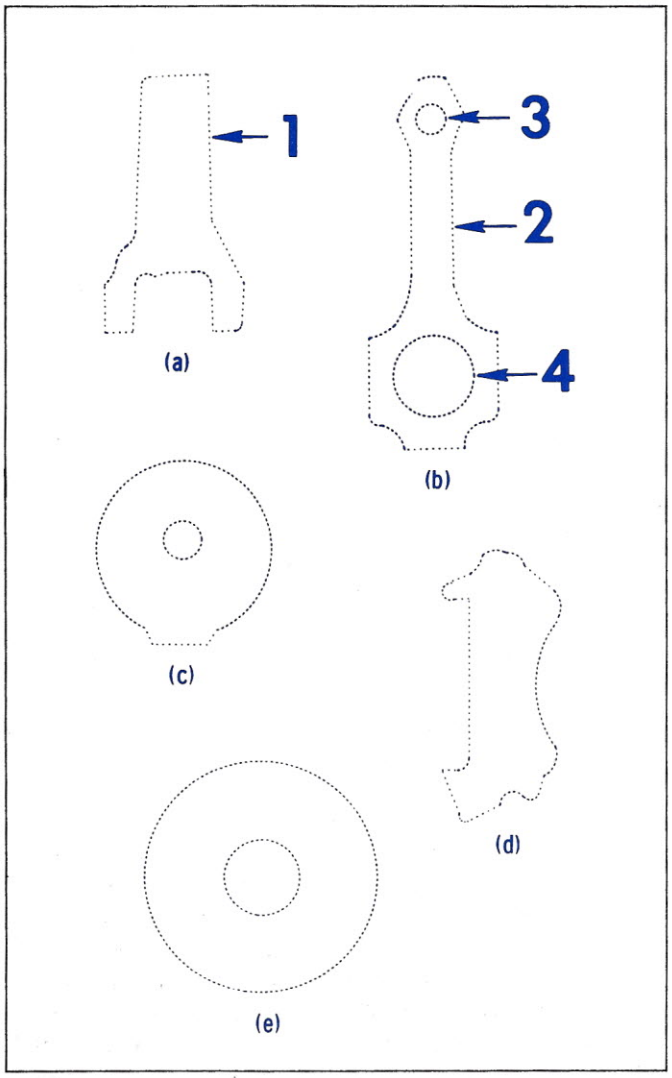

Figure 7. Five template models of automobile parts: (a) universal yoke, (b) connecting rod, (c) compressor body, (d) bracket, (e) gear blank.

including a scale transform in the comparison procedure, but this can add considerable computational overhead if the scale change is arbitrary.

Examples of applications meeting the requirements for successful use of template-matching include recognizing printed characters, the identification of navigation checkpoints by cruise missles, and, significantly, recognizing industrial parts on conveyer belts. In the last case, it would seem that template-matching would have trouble, since parts could be placed on the belt in different orientations, i.e., with a different side up than the one expected in the template. The rotation is usually not arbitrary, however, because most parts have a limited number of stable resting positions on the belt. To apply template-matching it is necessary to have a template for each possible stable position. If the orientation of the part around a vertical axis is not controlled, it will still be necessary to perform the comparisons with the templates either by a 2-D rotation of the template or by basing the comparison on rotationally invariant properties of the templates.

To illustrate these techniques, let's examine a system for parts recognition developed at the General Motors Research Laboratories [3]. In this system the templates are represented as sets of connected curves (straight lines and circular arcs) called "concurves". Each set of concurves forms a two-dimensional model of the outline of some known object in a stable resting position. The outer boundary of an object can be represented by a single closed concurve, and each interior hole in the object requires a concurve to describe it. Figure 7 shows five models; the numbers indicate individual concurves in a model. Each concurve is represented as a set of orientation-independent properties such as length, number of arcs, total angular change of all arcs, etc. Orientation-dependent data is stored separately as a set of vectors perpendicular to the concurves at regularly spaced intervals along them.

To recognize parts in a scene the system analyzes the results of edge detection to form a set of concurves that most accurately fits the observed edge data (Figure 8). Since transforming the orientation-dependent model data into image coordinates is computationally expensive, the system first compares the more abstract concurve properties to compute estimates of the "likelihood" of a match before the detailed transformations are made. Each concurve in the image is compared to each concurve in each model to compute the average likelihood of finding each kind of object in the scene. This requires N (number of concurves in scene) times M (total number of different

PHOTOS COURTESY OF GM RESEARCH LABS

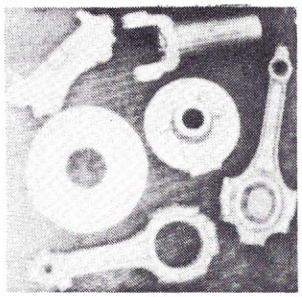

(a)

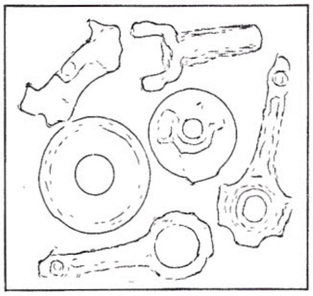

(b)

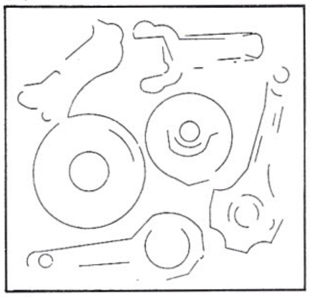

(c)

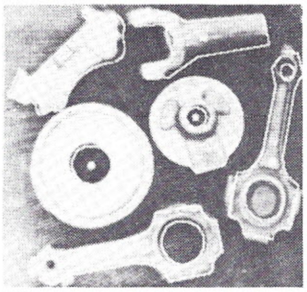

(d)

Figure 8. Recognition of parts on a conveyor belt: (a) Digitized scene, (b) Results of edge detection, (c) Concurves derived from edges, (d) Recognized models superimposed on the image.

concurves in the model database) comparisons, which can be large. However, this can be done fairly quickly since only the scalar properties are used. The vector description is transformed into the image only to verify the location of the most likely object, and to mark edges in the image as "used-up" so that they won't have to be considered in subsequent comparisons.

Because the actual templates are not used in the initial comparisons, this system is more accurately described as a form of "model-matching" based on the abstract properties of the concurves, but its restriction to two-dimensional patterns of a known size emphasizes its dependence on the templates. Also, since the template is used to verify a match and to mark edge data as "consumed", the system is able to work well even in the presence of visual noise (imperfect, fragmented edge data) and in cases of partial occlusion of the objects, as shown in Figure 9. This capability illustrates the power of template-related methods to deal with noisy images, since the template can tell the program what to look for and where to look for it. Systems such as this one could have wide application in manufacturing, both because of their ability to work well in the restricted environment and for the ease with which new objects may be added to its recognition repertoire. An example of a new object can be "shown" to the system under favorable conditions (high contrast, isolated object) and the program can compute and store both the template and its related properties. The speed of such systems (about 30 sec. for the complex scene in Figure 9) can be significantly increased by the use of special-purpose digital vision hardware such as that described above, making them suitable for the rapid interaction with robot manipulators required for assembly-line applications.

Recognition of Objects in a Three-Dimensional World

The GM recognition system illustrates how orientation-independent properties associated with the system's model of an object can be used to greatly reduce the need for template rotations when matching against edge data in the scene. Without such techniques, the matching process would be much more time-consuming. In 3-D environments, the additional transformations of scale, rotation, and perspective must be considered, and we immediately realize that the complexity added by each new degree of freedom makes the literal comparison of the image of some object with stored object models even more impractical. To compute the projection of a 3-D model of an "archetypical" object into an image is almost too costly to be used for anything other than verifying the system's most likely "guess" about what it is seeing.

Even the possibility of storing "literal" models of the surface structure of objects becomes impractical if the number of objects is large. The surface models used in computer graphics systems to describe the shape of curved objects sufficiently to compute a projection from

PHOTOS COURTESY OF GM RESEARCH LABS

Figure 9. A scene with occluded parts: (a) Digitized scene, (b) Derived concurves, (c) Recognized models superimposed.

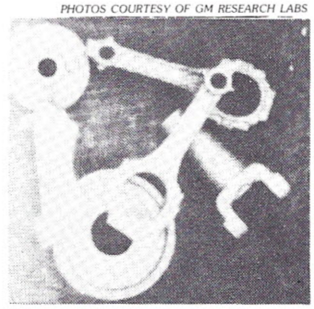

(a)

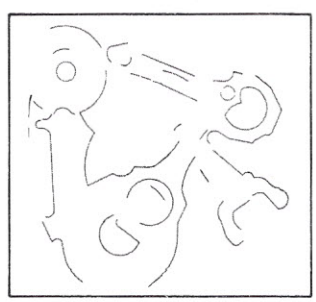

(b)

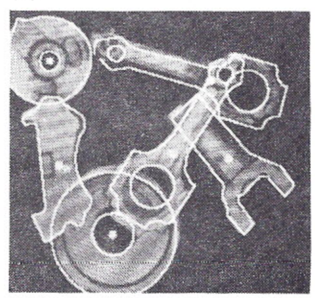

(c)

any perspective require thousands of bytes of storage. Since we do not know the capacity of human visual memory or the mechanisms behind its functioning, we cannot rule out the possibility that such literal surface models are somehow actually "stored", but merely the magnitude of the volume of information required makes the prospect unlikely. What seems more probable is that some form of encoding is used that compresses the unique characteristics of a particular kind of object into a highly condensed form.

Clues to the nature of the possible abstractions are suggested by examining drawings of familiar objects made by small children before they acquire the knowledge of perspective representation. These drawings often have the characteristic of showing structural relationships between the components of complex objects with little emphasis on the details of the components. (See Figure 10). The classic example is the stick figure representation of the human body. The fact that the child may be able to draw a more detailed representation of some individual part, such as an arm or a foot, suggests to some researchers that there may be a form of hierarchy in the understanding process. This may take the form of a pyramidal structure of definitions in which complex objects are described by successive layers of increasing detail, or in which recognition is accomplished by processes operating at different levels of detail. In each case, the top levels would be most concerned with the connections between the major components of the object and the lower levels with descriptions of primitive shapes. The ability of the human brain to "represent" and remember such structural knowledge would be intimately related to its aptitude for spatial relationships in general.

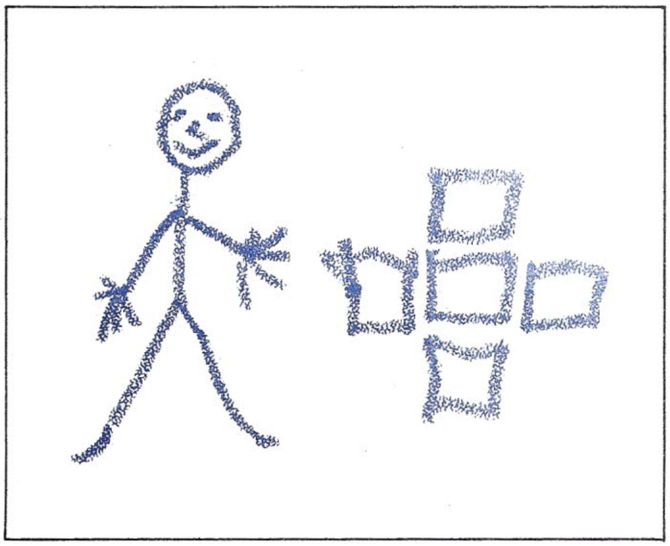

Figure 10. Drawings by children often show structural relationships without attention to actual appearances. The figure on the right is supposed to be a cube, and captures the significant feature of the squareness of the cube's faces.

To translate this neuro-physiological theory of vision into a computer program, a researcher must decide what kind of knowledge is needed to accomplish recognition and how it can be represented and used effectively. For systems in which recognition procedures operate on descriptions of the objects, this may involve specifying a vocabulary and a grammar for expressing structural relationships within a particular level of detail and also for describing the connections between levels. In the GM recognition system, the top level of the description of an object contains the derived abstract numeric properties of the templates, with the description of its actual shape left to the lower level. In models of three-dimensional objects, the abstract structural relationships describing the object at the higher levels may have a fairly straightforward translation into English, using relational terms such as "Is-a-part-of", "Is-a-kind-of", "Is-connected-to", "Consists-of", etc., applied to terms denoting parts of the object at various levels of the description.

Consider the child's stick figure drawing of the human body. Such a drawing corresponds closely to a description of the connectivity of the limbs, and forms an appropriate top level description. The features of such a drawing can be translated into a computer data structure using relational terms, forming the description represented in Figure 11. From each named "node" in the description, other connectives would point to the description of that component in the next level of detail. These relations form a directed graph structure, in which the description of an individual object would form a "tree" with the name of the object as the top node. Note that the top-level description need not refer to any numeric properties of the image of the object, such as the shape of the limbs or the angle at which they join to the torso. Since the object may appear in a picture viewed from any perspective or with the limbs held at any angle (within their range of movement), such details would not be very useful in recognition. The abstract description captures the essential features that should be invariant in any image of the object.

Of course, at some level of the description numeric detail is essential. Many kinds of objects may have similar connectivity descriptions at the top level. For example, the connectivity graph for the nodes in the top level

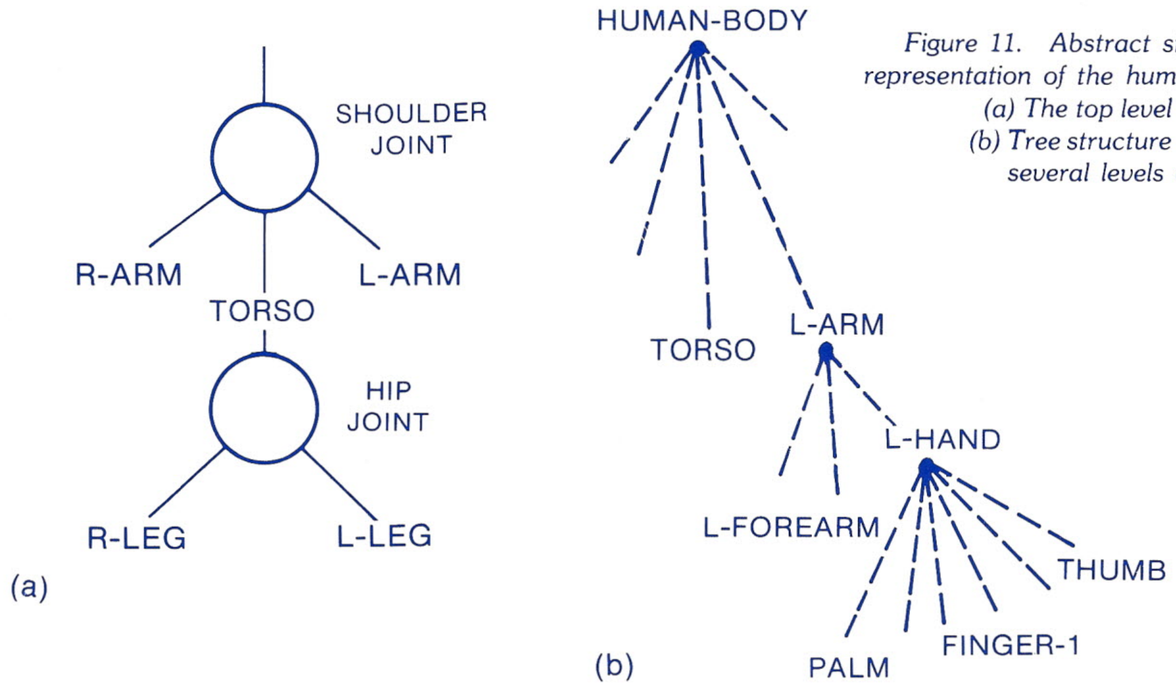

Figure 11. Abstract structural representation of the human body: (a) The top level of detail, (b) Tree structure showing several levels of detail.

description of the human body applies equally well to the description of many mammalian bodies. The only way to discriminate between each kind of mammal is to examine the sizes and shapes of the components. Comparison at this level of detail is computationally expensive, and should only be performed after descriptions of potential matches are selected at the top level, similar to the procedure used in the GM system.

The method of representing the shape descriptions is a major design consideration. The large amount of storage required for numeric surface models and the difficulty of comparing them suggests a need for a compressed representation that may be easily computed from an image and easily compared with the shape descriptions in the stored models. This is especially true for the curved surfaces typical of objects found in nature. For manufactured objects, modelling by geometric solids may be sufficiently concise and quite accurate.

One form of shape description that can work well for curved objects is the "generalized cone" representation introduced by Binford [4]. In this scheme, the shape of a component may be described by specifying an axis and successive cross sections of the object in planes perpendicular to the axis (Figure 12). For many objects, the choice of an axis may not be unique, but for elongated objects the obvious choice is one along the longest dimension that passes through the centers of the successive cross sections. The axis itself may be a curve, possibly represented as a series of line segments between the cross sections. The cross sections should be single closed curves; if the object branches, a joint should be declared and new axes specified for the separate branches. This representation works well for objects whose cross section varies smoothly along the axis. Further data compression can be obtained by storing statistics and properties derived from the cross sections, such as the average area, its ratio to the length of the axis, whether the cross sections are more cylindrical or conical,

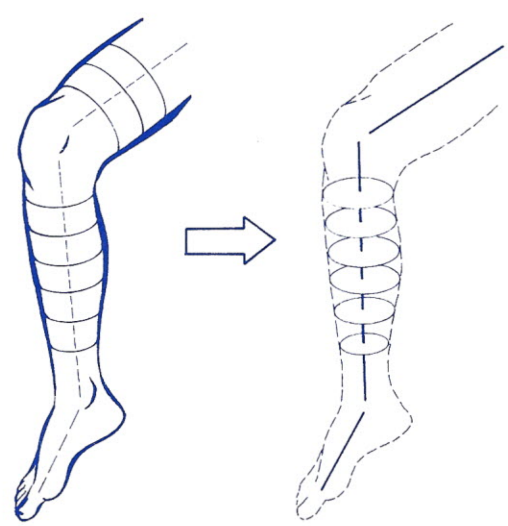

Figure 12. The "Generalized Cone" representation of curved solids.

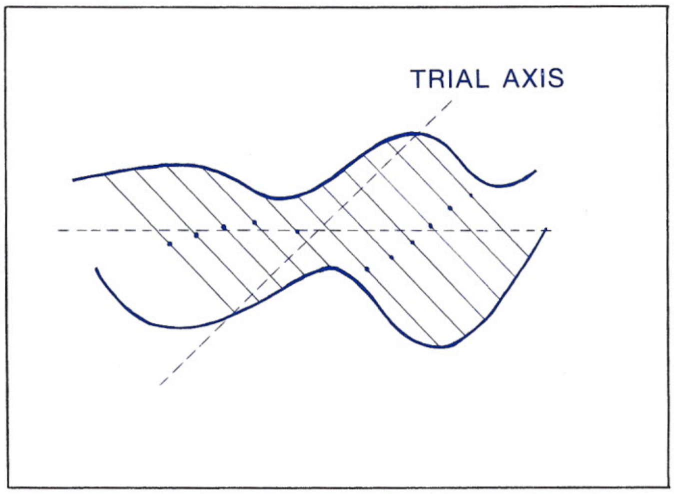

Figure 13. Finding the best choice for the axis of a region by iteration.

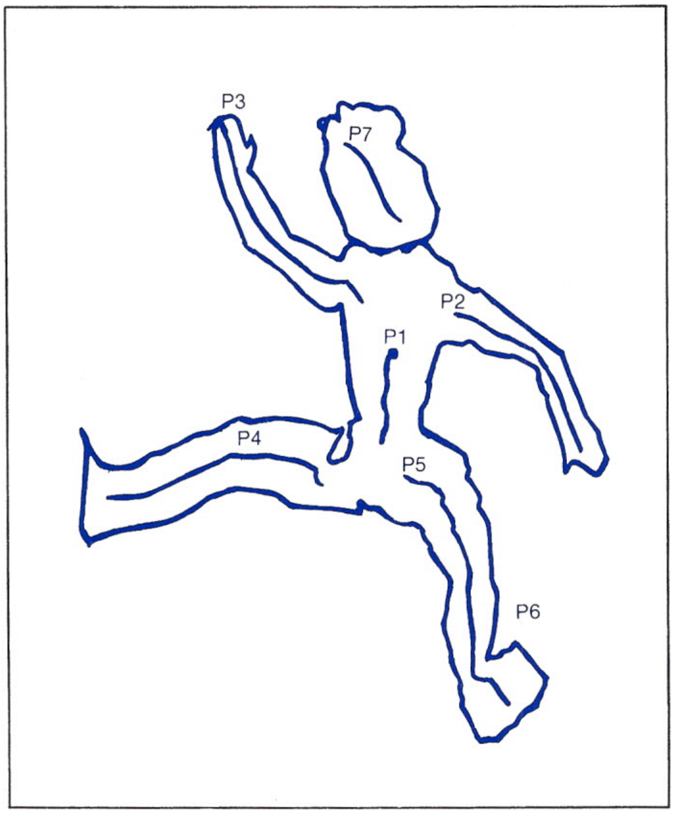

Figure 14. Cones computed from the edge picture of a toy doll.

etc., instead of storing the actual plane curves for each cross section along the axis.

Although restricted to two-dimensional recognition, the GM system demonstrates some of the characteristics essential for systems operating in the unrestricted 3-D world. To reduce the need for time-consuming coordinate transforms and literal image comparisons, a description of the image is formed in terms of the available vocabulary of image primitives and relational connectives. Properties derived from the description are used as an index into the database of stored models and likely candidates are selected. Then, discrimination between the candidates is based on comparisons at levels of increased detail. In more elaborate systems, the program can use the more detailed levels of the model to tell it where to look for particular features in the image that may not have been in the first-pass description of the observed object. Similarly, a statistically good match between some region in the image and a component of some known object can be used to direct the search for the other components of that object in the image.

A working recognition system built on the concept of model descriptions of objects in terms of generalized cones has been described by Nevatia and Binford [5]. Although the image input to the system is from a laser rangefinder that returns the three-dimensional coordinates of object surface points, the same techniques could be applied to 2-D images, though discrimination based on the absolute size of parts of objects would not be possible (unless stereo vision input is available). Edge detection in a laser range image involves filtering the image array for discontinuities in range, rather than in contrast, and results in an analogous edge picture in which the 3-D location of the edge points is also known. Features of the top level description of an object in an image are used to index into the model database, selecting model descriptions for more detailed comparison.

The system works by first computing a generalized cone description of all the objects in the scene, using the edge picture. Since the axes of parts of the objects are not initially known, producing the description is an iterative process in which axes in several directions are proposed. For each trial axis, 2-D "cross-sections" of the object are constructed that are normal to the axis and at regularly spaced intervals along it (Figure 13). By fitting a new axis to the midpoints of the cross sections, a better choice may be found for the next iteration. If the process does not converge in a few steps, then the region may have no well-defined axis. If a good choice of axis is found, the axis is extrapolated at either end to cover as much as possible of the region. The extrapolation is terminated if the end of the region is encountered or if a discontinuous jump in the size of the cross section indicates that a joint with some different region may have been reached. Figure 14 shows the result of computing generalized cones to describe the image of a toy doll.

The next step in the recognition process is to construct a symbolic description of the generalized cones for each

object. Regions described by a single axis are characterized as conical or cylindrical, and shape statistics are computed. The areas between well-defined axes are identified as joints, and may be described according to the number of axes they connect and the spatial configuration of the connected axes. The result is a connectivity graph (Figure 15) which (hopefully) is similar to the graph for the stored model of that kind of object. If there are disconnected regions adjacent to some object in the scene, potential joints are hypothesized that connect them to the object and may be verified or rejected during the recognition process.

Finally, the resulting description must be matched against the stored models. Since the computed descriptions may differ from the models due to imperfect edge evidence, missing parts, extra parts, distortion due to the particular perspective in the scene, etc., this matching must be statistical as well as literal. A key problem is finding the smallest possible number of candidates to compare. The solution taken in this system is to use features of the object description to compute an index into the model database. Similar indices are computed for each stored model, so that given an index computed from an object in the image, a list of models with the same index is immediately available. Several such indices may be computed for a single model, based on features of the largest components in the model. Once the smallest set of candidates is found, the actual matching of descriptions can take place. The match with the best score (based on the statistical matches of the individual pieces) is then verified to determine if differences between the object and model descriptions can be satisfactorily explained. In this case, the object description of Fig. 16(b) matched the graph of Fig. 12(a) better than the other candidates in the program's database, despite the extraneous joints J4 and J5 observed in the image.

If pieces of a candidate model are not matched in the object model, potential joints identified in the description phase may be examined to locate the missing pieces. Figure 16 shows another view of a toy doll in which, due to image noise, the regions for an arm and a leg are disconnected. The matching process correctly joined the two regions with the larger description and concluded that the object was a doll, the next most likely candidate being a toy horse, whose connectivity graph is the same as the doll's except for the addition of a tail. The match with the toy horse was rejected based on the differences in the statistical matches of the individual pieces. (Inclusion in the model of the possible range of movement of the joints could have helped rule out the match with the horse as

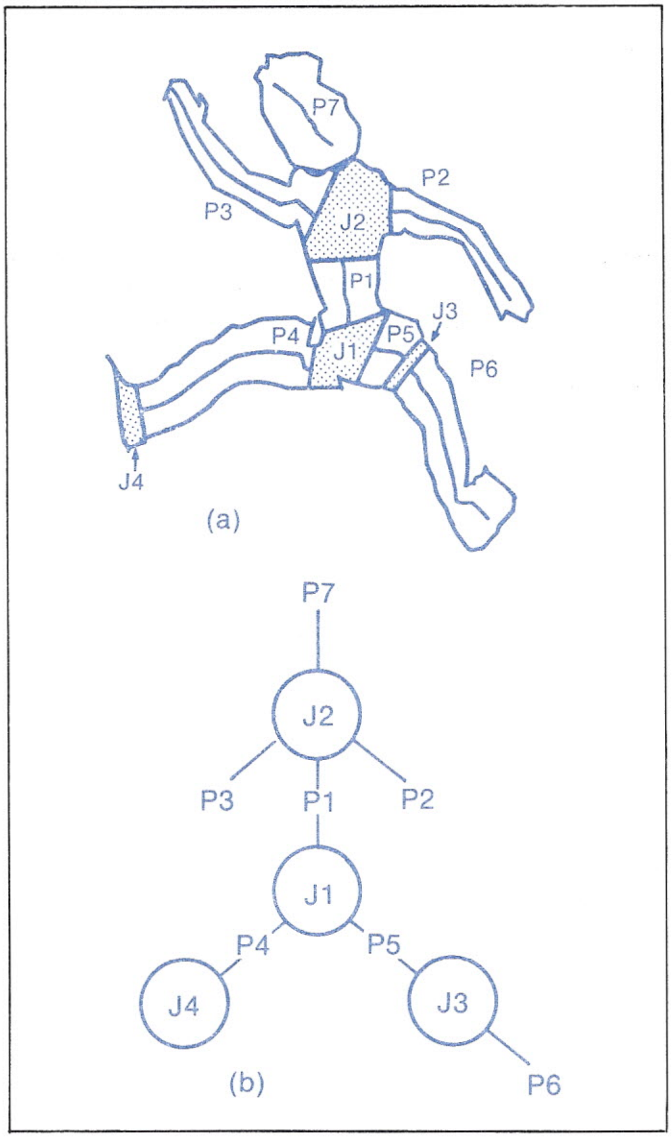

Figure 15. Resulting top-level description of the doll: (a) Cones connected by joint regions, (b) Resulting connectivity graph. Note the extra joints (J3 & J4) derived from the cones.

well.)

This system requires several minutes of computer time (on a DEC KA-10 computer) to process an image and recognize the objects, even with only five stored models. Most of the time is spent computing the generalized cone description from the edge data, a process which, like the description phase of the GM recognition system, can be greatly speeded up by the use of hardware edge detection and simple parallel processing schemes. The effective use of parallel processing may ultimately be the key to fast vision systems. The processes of description generation, model comparison, and verification all involve repetitive application of standard procedures to different sets of independent image and model data. If these procedures could be applied in parallel, the results could be made available to the next step in the recognition process

immediately upon completion (with suitable process synchronization). This would result in a speed increase roughly proportional to the degree of parallelism.

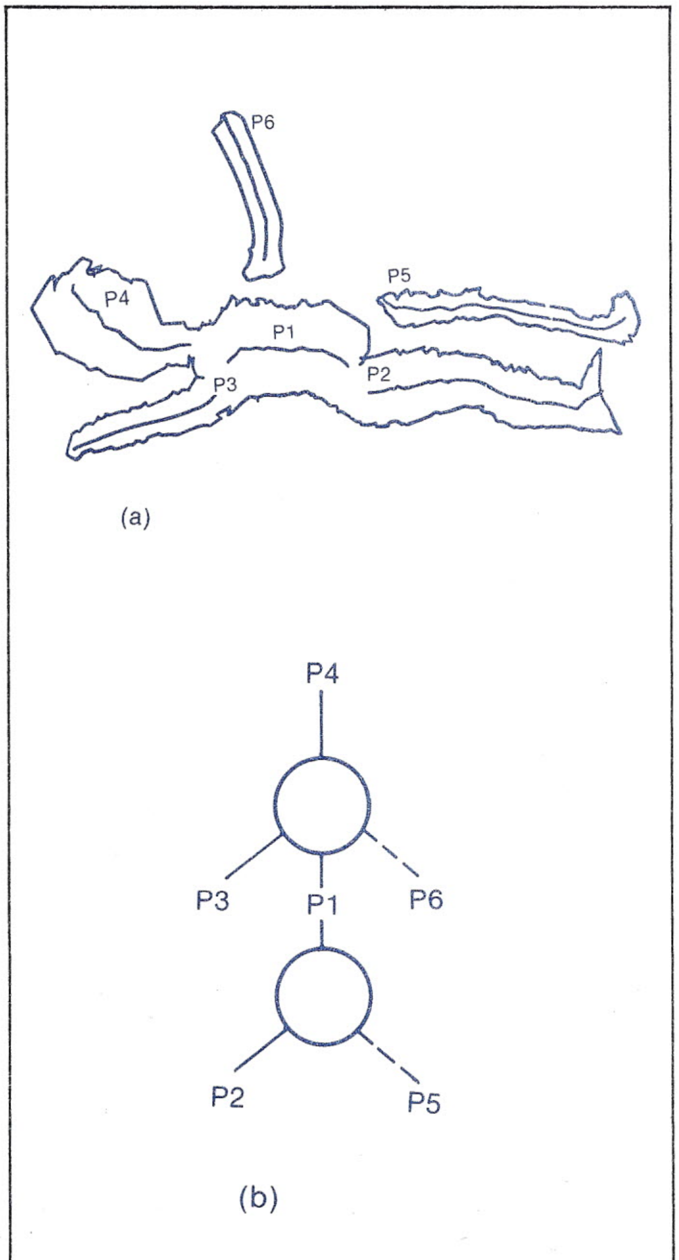

Figure 16. Another view of the toy doll with increased image noise and a different orientation of the limbs and body. (a) Edge data and derived cones, (b) Resulting connection graph. A potential joint connects the left arm and leg to the rest of the graph.

Summary

This article has attempted to give the reader some insight into the complexity of machine recognition. The intent was not to provide an exhaustive survey of the field of automatic perception, but rather to concentrate on examples that illustrate the variety of processing involved and an indication of the capabilities of state-of-the-art systems. Discussion of the many successful systems dealing with restricted domains, such as those limited to the recognition of toy building blocks and other geometric solids, or image segmentation based on texture descriptions, or statistical pattern classification, is beyond the scope of any single article. Hopefully, by considering some representative systems, the character and methodology of computer vision research may be conveyed.

The processes that underly image understanding extend beyond the problem of recognition. Indeed, the challenge of representing knowledge and its effective application to solving problems underlies all of automatic cognitive inference and is a central theme of Artificial Intelligence. Computer vision in particular offers immediate benefits for increased manufacturing productivity as well as many other areas of automation. Research in this area, as well as in other aspects of cognition, will almost certainly have a profound impact on the use of computers in industry, space exploration, and domestic applications.

References:

[1] R. Duda and P. Hart, *Pattern Classification and Scene Analysis,* New York, Wiley and Sons. (1973).

[2] R. Eskenazi and J. Wilf,"Low-level Processing for Real-Time Image Analysis", IEEE Workshop on Computer Analysis of Time-Varying Imagery, Philadelphia, Pa. (April 1979).

[3] W. Perkins, "Model-Based Vision System for Scenes Containing Multiple Parts", Proc. of the Fifth International Joint Conference on Artificial Intelligence (IJCAI5), Cambridge, Masschusetts, (Aug. 1977).

[4] T. Binford, "Visual Perception by a Computer", IEEE Conf. on Systems and Controls, Miami (Dec. 1971).

[5] R. Nevatia and T. Binford, "Description and Recognition of Curved Objects", Artificial Intelligence 8, (1977), North-Holland Publishing Co.

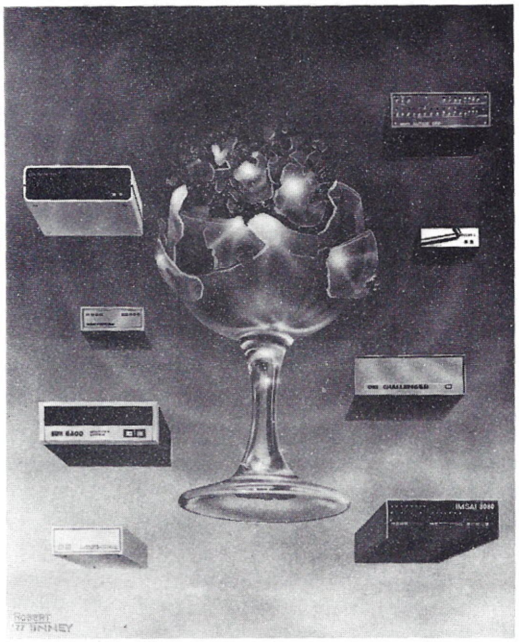

250/750 BREAKING THE SOUND BARRIER Robert Tinney

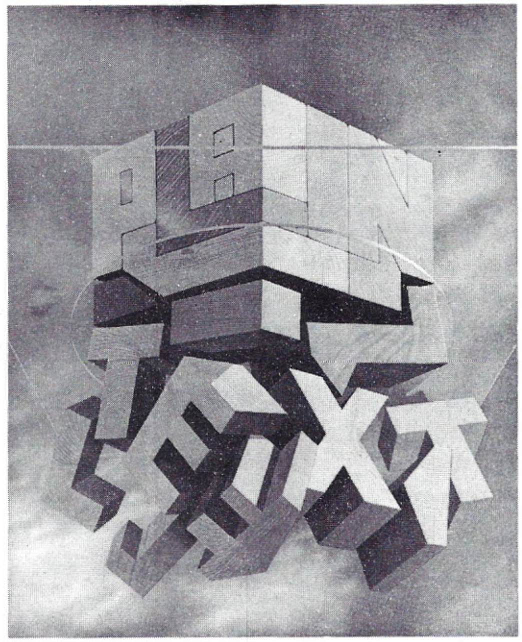

250/750 THE TRAP DOOR Robert Tinney

March 1979

Byte Cover Prints -- Limited Editions.

The September '77 and March '79 covers of BYTE are now each available as a limited edition art print, personally signed and numbered by the artist, Robert Tinney.

These prints are strictly limited to a quantity of 750 for each cover, and no other editions, of any size, will ever be published. Each print is 18" x 22", printed on quality, coated stock, and signed and numbered in pencil at bottom.

The price of each print is $25. This includes 1) a signed and numbered print; 2) a Certificate of Authenticity, also signed personally by the artist and witnessed, attesting to the number of the edition (750), and the destruction of the printing plates; and 3) first class shipment in a heavy-duty mailing tube.

To order your limited edition art print, fill out and mail the order form below.

Circle 20

Send me ________ "Breaking the Sound Barrier" prints at $25 each, and ________ "Trap Door" prints at $25 each. I understand this price includes Certificate of Authenticity and first class shipment.

☐ I have enclosed check or money order to Robert Tinney Graphics.

☐ Charge this to my Master Charge or Visa

Card #________________ Expires:______

Ship my print(s) to:

Name________________________________

Address______________________________

City________________ State______ Zip________

Send order to:

robert tinney graphics

P.O. Box 45047 • Baton Rouge, LA 70895

The Grivet CHESS PLAYING ARM

PROBLEM: How to position a robot arm in three dimensions without the time-consuming floating point trigonometric calculations required for angular joints? How do you eliminate the need for analog to digital converters and still have accurate position feedback for precision control? How do you minimize the amount of external circuitry required for the computer to drive the arm?

The solution to each of these problems is found in the design of the GRIVET Series-2 robot manipulator. GRIVET stands for "Gallaher Research Inc. Versatile Electromechanical Tool", after the company started by its inventor, John Gallaher (who tells us that it is also the name of a species of small monkey). The simplicity of the GRIVET arm design has resulted in a mechanism ideally suited for interfacing to and control by microcomputers, but with performance comparable to some expensive industrial manipulators.

Most industrial robots have arms composed of two or more segments connected by angular joints (although many have linear linkages as well.) The joint angles of these arms are typically controlled by programmed analog servos and/or fast minicomputers whose floating point arithmetic hardware can handle the volume of trig computations needed to position the arm in the workspace. The GRIVET/3 arm, however, positions its hand in the workspace by the linear movement of each of its three linkages, so that the hand coordinates may be computed from a simple linear translation of the position of each link along its range of travel. A manipulator using this linkage design is referred to as a "Cartesian" arm, since the translation directions of its linkages define Cartesian coordinate axes for its workspace and because of the direct correspondence between link position and workspace coordinates. Of course, there is a price for this simplicity, and that is that a Cartesian arm must necessarily be at least as big as the workspace it is capable of covering. The GRIVET/3 arm is built so that the links can travel 18″ in the X or Y directions and 6″ vertically, easily enough to cover a standard chessboard!

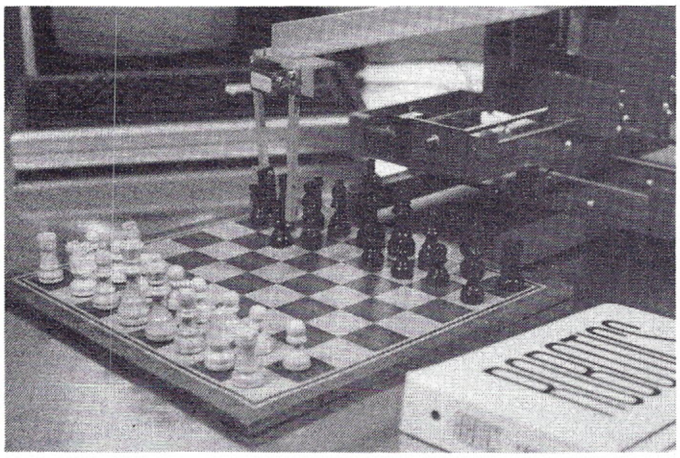

Although the linear linkages remove the burden of angular joint calculations, there still remains the problem of accurately determining the position of each link along its axis. Again, the solutions to this problem found in most industrial manipulators tend to be expensive. Most such robot arms feed their joint positions back to their controlling computers either by coupling a potentiometer (a variable resistor turned by a shaft) to each of the angular joint shafts, or by using an optical position encoder attached to the joint shaft to produce a digitally encoded measure of the joint angle. Arms using a "pot" wired as a voltage divider to sense the joint angle require an analog to digital converter (ADC) with one input channel for each joint. Optical shaft encoders are expensive, and multichannel ADC's also get expensive depending upon the accuracy desired, as well as adding to the external

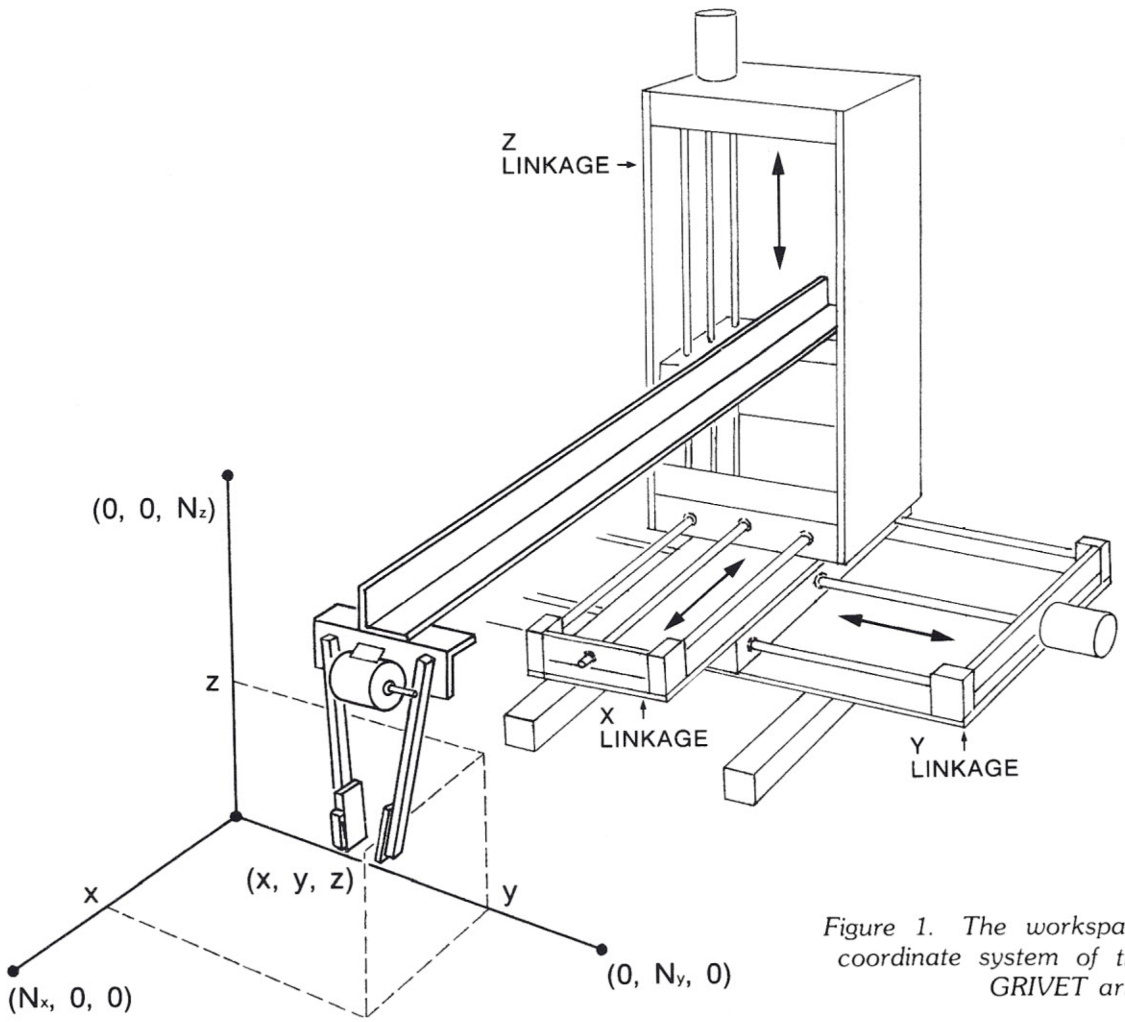

Figure 1. The workspace coordinate system of the GRIVET arm.

circuitry required by the system. Another problem with using pots for position sensing is that expensive heavy duty pots must be used to handle the stress of continual turning. The problem of accurate positioning can also be solved by using stepping motors. These are multiple pole DC motors that can be "pulsed" by drive circuitry to cause the motor shaft to "step" from one stable pole position to the next. Stepping motors are considerably more expensive than ordinary DC motors of comparable torque, and in addition, they usually require involved drive circuitry to produce the power pulses.*

*A reduction to the drive circuitry overhead for stepping motors is made possible by Philips, the British electronics firm, some of whose steppers are completely interfaced to TTL with one 16 pin DIP, needing only a direction bit and a TTL pulse input to step the motor.

Here again, the GRIVET/3 arm is designed with simplicity and economy in mind. Since each of the three links is positioned along its range of movement by a lead screw turned by a motor (see mechanical discussion below), the number of turns of the screw shaft can be used to give an accurate measure of the movement of the link with respect to a known position. Gallaher's approach is to use inexpensive optical sensors to determine a "home" position for each link and to allow the controlling computer to count the turns of the screw shaft. If the computer program controlling the arm measures workspace coordinates in terms of shaft turns from the home position, no arithmetic is required to position the arm to a desired location. (See Figure 1.) In addition, the use of a simple solid state relay to switch each motor's drive current takes the place of a linear power output circuit while still offering sufficient speed control for most applications.
Here's how it works...

Mechanical Description

The basic linear linkage of the GRIVET/3 arm is composed of a carriage that slides along a pair of stainless steel guide shafts (Figure 2). Linear ball bearing bushings are used in the guide holes in each half of the carriage, and the separation of the halves reduces off-axis play to a minimum. The basic structural material is anodized aluminum. The carriage is positioned along its range by a motor-driven lead screw, similar to the head positioning mechanism of most floppy disks, machine tools, etc. In this case, the screw pitch is .05", so that if the controller counts the turns of the lead screw, the position of the carriage is known to within this tolerance. The motor used is an Igarashi Electric Works 5-9 Volt model, capable of supplying approximately 4 in.-oz. torque at 12,600 rpm, drawing 8.8 amps. The GRIVET specifications lists transverse speeds of 10"/sec. indicating a maximum lead screw speed of 12,000 rpm. The carriage subassembly uses threaded inserts to couple to the lead screw, and the free end of the lead screw is held by a nylon bearing in the end plate. At each end of the frame is a microswitch which will be tripped by the bottom of the carriage if it comes too close to the end of its range of travel. The switches can be wired in series with the motor to cut the power should an irresponsible or buggy program try to run the carriage off the end. (More adventuresome builders can use them merely to warn the controller or to activate an override.)

Each of the three arm axes is defined by such a linear linkage, starting with the Y axis which rests on the base and provides the sideways movement of the arm. Next, the X axis linkage is built on the carriage of the Y axis link, rotated 90 degrees to provide the forward-backward movement. The Z linkage is built vertically on the X carriage. Although the photograph of the prototype shows only a single carriage subassembly, the production model uses the same dual carriage assembly as on the X and Y links to minimize off-axis play. A 22" boom is attached to the Z axis carriage, reaching out over the workspace. The

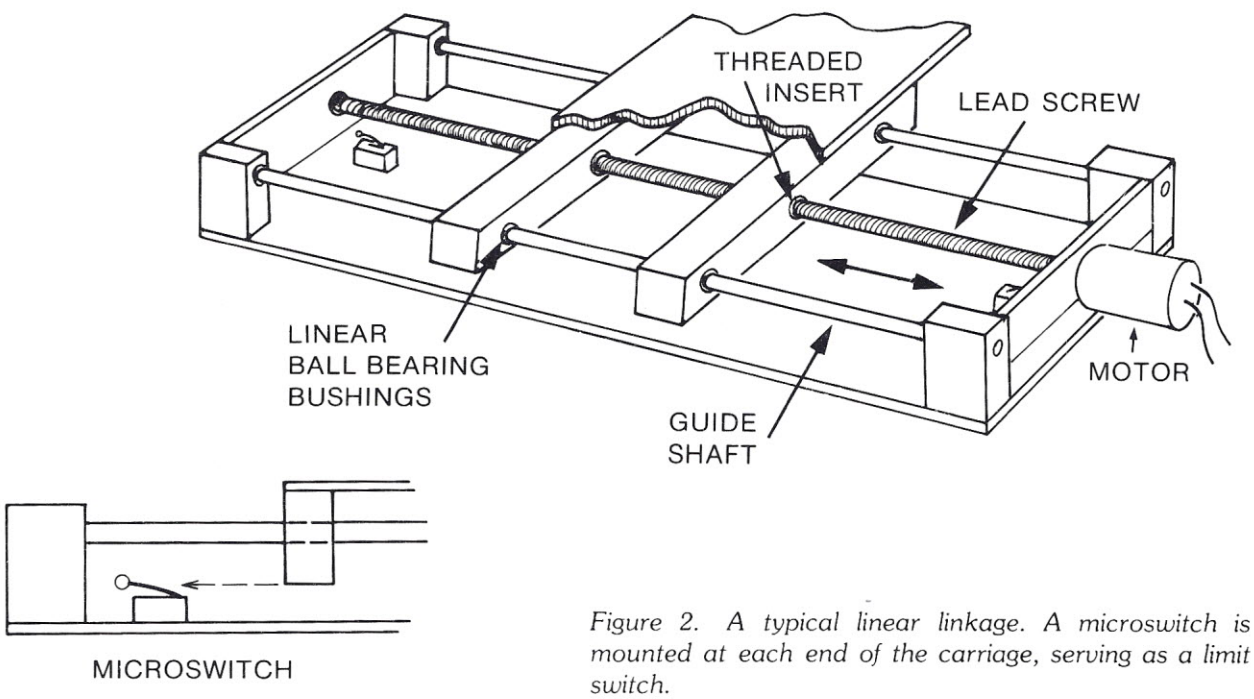

Figure 2. A typical linear linkage. A microswitch is mounted at each end of the carriage, serving as a limit switch.

Figure 3. The end-effector. The microswitch on the stationary finger serves as a grasp sensor.

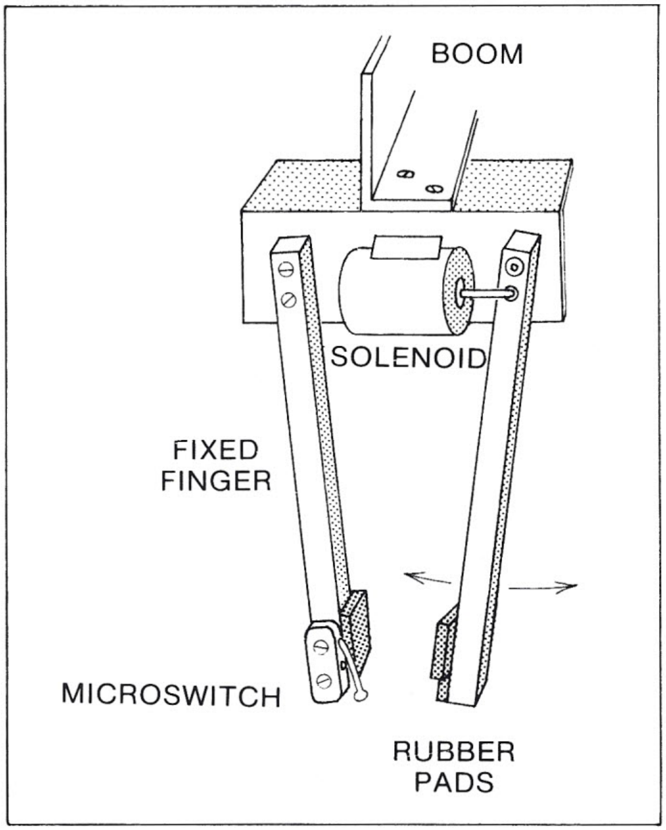

end-effector, a two-fingered gripper, is affixed to the end of the boom reaching down. One of the fingers is stationary, and the other is controlled by a solenoid, as shown in Figure 3. Rubber pads on the fingertips help to prevent slippage when grasping. A microswitch "touch sensor" is mounted in a niche in the rubber pad on the fixed finger so that it will be tripped when an object is grasped between the fingers. If the fingers are closed and touching with no object between, the switch lever fits into a corresponding niche in the opposite pad, so that the switch remains open. This allows the robot to search for objects by successively closing, opening, and moving the hand until a grasp signal is obtained.

Measuring the Arm's Position

An indispensible requirement of all robot systems is for the controlling computer to have accurate knowledge of the position of its appendages. In the absence of visual feedback, (which is a difficult problem in itself,) position sensors built into the mechanical linkages must be used. To avoid the problems and expense involved with analog feedback or shaft encoders, the GRIVET/3 uses digital feedback from photoelectric sensors that count the turns of the lead screw on each joint. Each sensing element is a commercially available package, shown in Figure 4, consisting of an LED paired with a phototransistor. When wired and mounted as shown, the voltage at the collector of the phototransistor will vary inversely with the amount of light reflected by the surface. With sufficient contrast between the light and dark areas of the surface being sensed, the output of the sensor circuit can drive a TTL input directly.

The GRIVET uses two of these sensors on each linkage, mounted around the shaft coupler that connects the motor to the lead screw (Figure 5). Half of the circumference of the coupler is painted flat black to produce a TTL "high", and the reflectivity of the natural

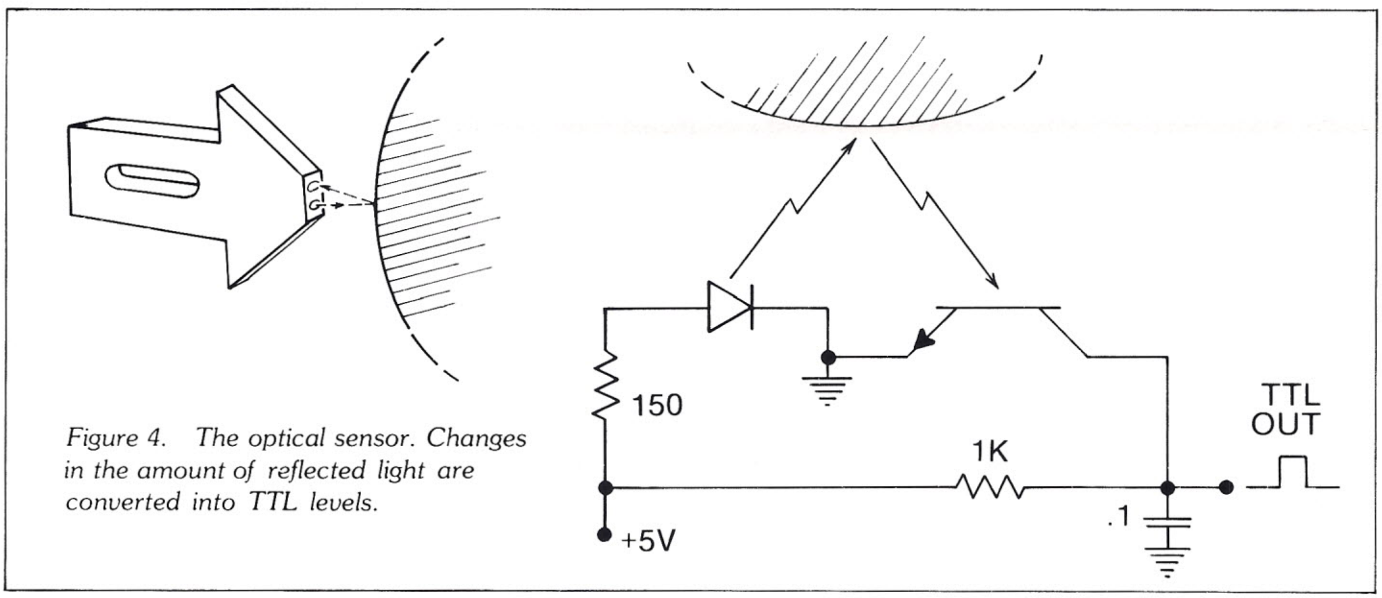

Figure 4. The optical sensor. Changes in the amount of reflected light are converted into TTL levels.

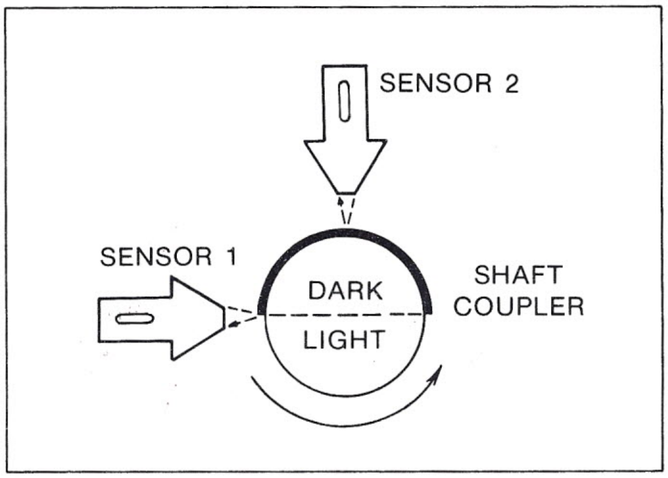

Figure 5. End view of the shaft coupler, showing the detectors 90° apart, allowing quarter turn resolution.

metal is sufficient to produce a "low". By interfacing the signal to the computer's interrupt line, (as well as to an input bit so the signal can be read) the computer will get an interrupt for each revolution of the shaft. The second sensor is mounted 90 degrees away, so that by reading its level the computer can tell the direction of rotation. If the interrupt line is connected so that the interrupt occurs when sensor 1 sees the dark side of the shaft coupler, then if sensor 2 is also dark, the rotation is counter-clockwise (as shown). Given the four possible combinations of each sensor seeing either light or dark, the computer can determine the shaft angle within 90 degrees. However, due to the "backlash" in the lead screw to carriage coupling, etc., this does not result in 1/4 turn accuracy in the carriage position. Play in the system linkages limits knowledge of the hand's position to about .1". At the top carriage speed of 10 in./sec., the 12,000 rpm shaft speed will produce 200 interrupts per second. Even with all three links at top speed, this still requires less of the computer's time than it takes to output to an average (9600 baud) CRT display terminal.

The interrupts from the sensors allow the computer to increment or decrement a word containing the carriage position measured in turns. It is still necessary to establish a reference position on the linkage to which the computer can send the arm and initialize the position counter when the robot is first turned on or whenever a reset is necessary. For this purpose, each link has an optical interrupter that is activated by a thin metal blade mounted on the bottom of the carriage (Figure 6). To locate the reset position, the controller drives the carriage toward the sensor until the leading edge of the blade is sensed. It can then "home in" to find the exact turn where the blade first cuts the optical path, and initialize (normally to zero) the position counter. The locations of the reset sensors on each axis together define the "home" position of the arm. They are normally placed so that the home position is the origin of the workspace coordinate system and only positive workspace coordinates are needed, as in Figure 1.

The Motor Control Circuit

Depending on the kind of application and the accuracy of control desired, there are many alternatives for regulating the power of a DC motor by a computer, each with its advantages and disadvantages. Linear output amplifiers work well for servo circuits of realtively low power, but power dissipation in the output transistors can become significant when driving motors of more than a few amps. The various forms of regulation by power switching are far more efficient, but usually require more control circuitry to achieve the same accuracy. When a computer is used as a motor controller, however, many of the functions of the switching regulator control circuitry can be performed in software, reducing external circuitry requirements. This is the case in the power circuit used by the GRIVET/3 arm. By performing the control functions in software and by relaxing the requirements for precise velocity control of the arm linkages, motor control is accomplished with a minimum of external circuitry.

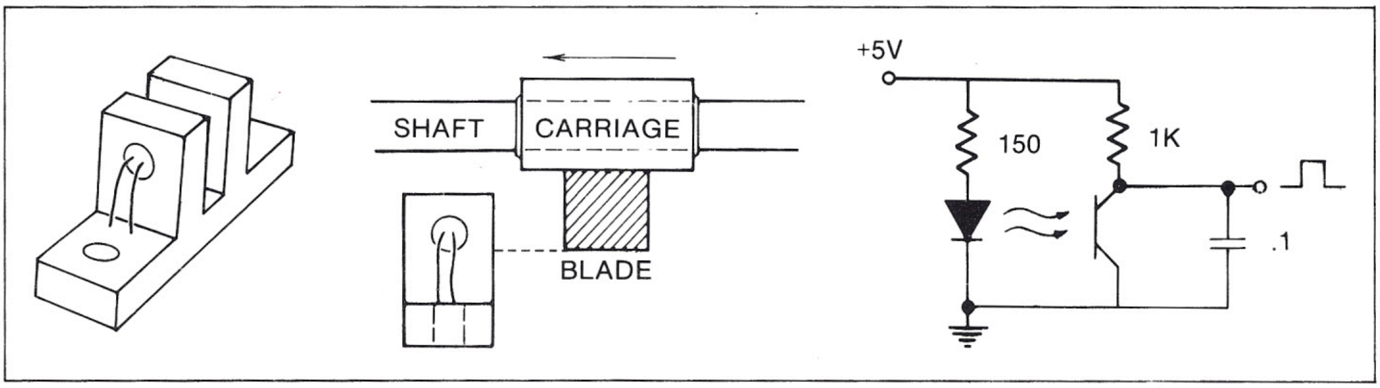

Figure 6. An optical interrupter, similar to the reflectivity sensor, is used to detect the "home" position.

The basic control circuit is shown in Figure 7. The only power supply is a line voltage transformer producing 12.6 Volts AC! The rectification necessary to drive a DC motor is accomplished by computer software selectively switching a solid-state relay (SSR) module wired in series with the motor. The zero-crossing detector converts the 60Hz input signal into a 60Hz TTL square wave. The computer interface will cause an interrupt to occur on either the rising edge or the falling edge. Depending on the motor direction desired, the computer can send a turn-on pulse to the SSR at the appropriate edge of the square wave, which will cause one-half cycle of the desired polarity to be delivered to the motor. (Leaving the turn-on signal low longer than half a cycle or sending pulses during cycles of different polarities would cause AC to be given to the motor, with consequent ill effects!) To accomplish speed control, the computer merely skips some of the possible cycles, and the motor averages out the half-wave rectified pulses to a lower voltage and thus runs slower. In fact, DC motors run better at low speeds on pulsed DC rather than a constant level of the same average voltage, because the pulse of torque helps to minimize the effects of friction that become more significant at low speeds. Since the pulse of power delivered to the motor is half a sine wave, the turn-off at the end of the cycle is smooth, greatly reducing the noise spikes usually found in square-wave switchers which have a way of spreading throughout the system.

Again, using the computer to provide the power synchronization adds to the overhead required for servicing the interrupts, but even with the additional 120 interrupts/sec. from the zero crossing detector, the maximum rate for controlling the entire arm comes only to roughly 720/sec., still hardly any burden to the controller. Also, the same control technique can easily be adapted for DC use in self-contained robots by providing a reference clock so the computer can measure the switching cycle of an output transistor. In this case, however, a reversing circuit or relay will be needed to change the motor direction.

Control Program Design

Because of the asynchronous nature of this interrupt-driven control method, the computer program that uses the arm should be structured as a set of independent modules that communicate by setting and reading common command and feedback variables, as shown in Figure 8. The chess-playing program issues commands in normal chess notation to its command translator, which converts the command into the coordinates of the piece to be moved and sets the command word for the arm control loop. From that point on, control of the CPU is determined by the actions of the I/O interrupts and a scheduling module. The position feedback module for a given axis is run whenever the shaft sensor interrupt occurs for that link, and increments or decrements the position counter for that link. The other modules are run at periodic intervals determined by the 120Hz power phase interrupt and the module's "priority", which indicates the urgency of its need to run. The following is a brief discussion of the "real-time" programming method normally used to control concurrent processes in a single CPU.

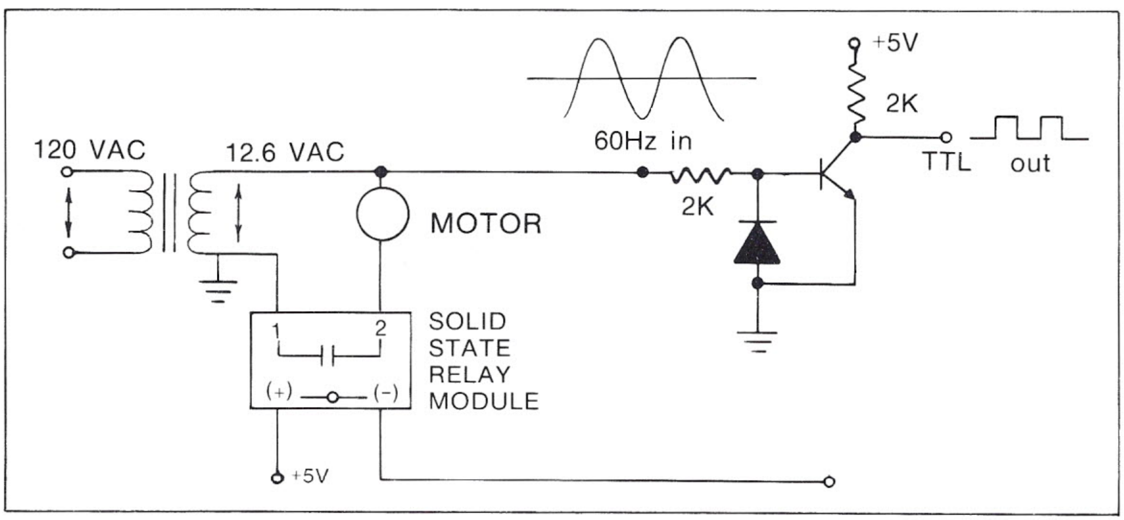

Figure 7. The motor drive circuit. A solid-state relay is selectively pulsed to rectify the AC. A zero-crossing detector produces TTL pulses that serve as a system clock.

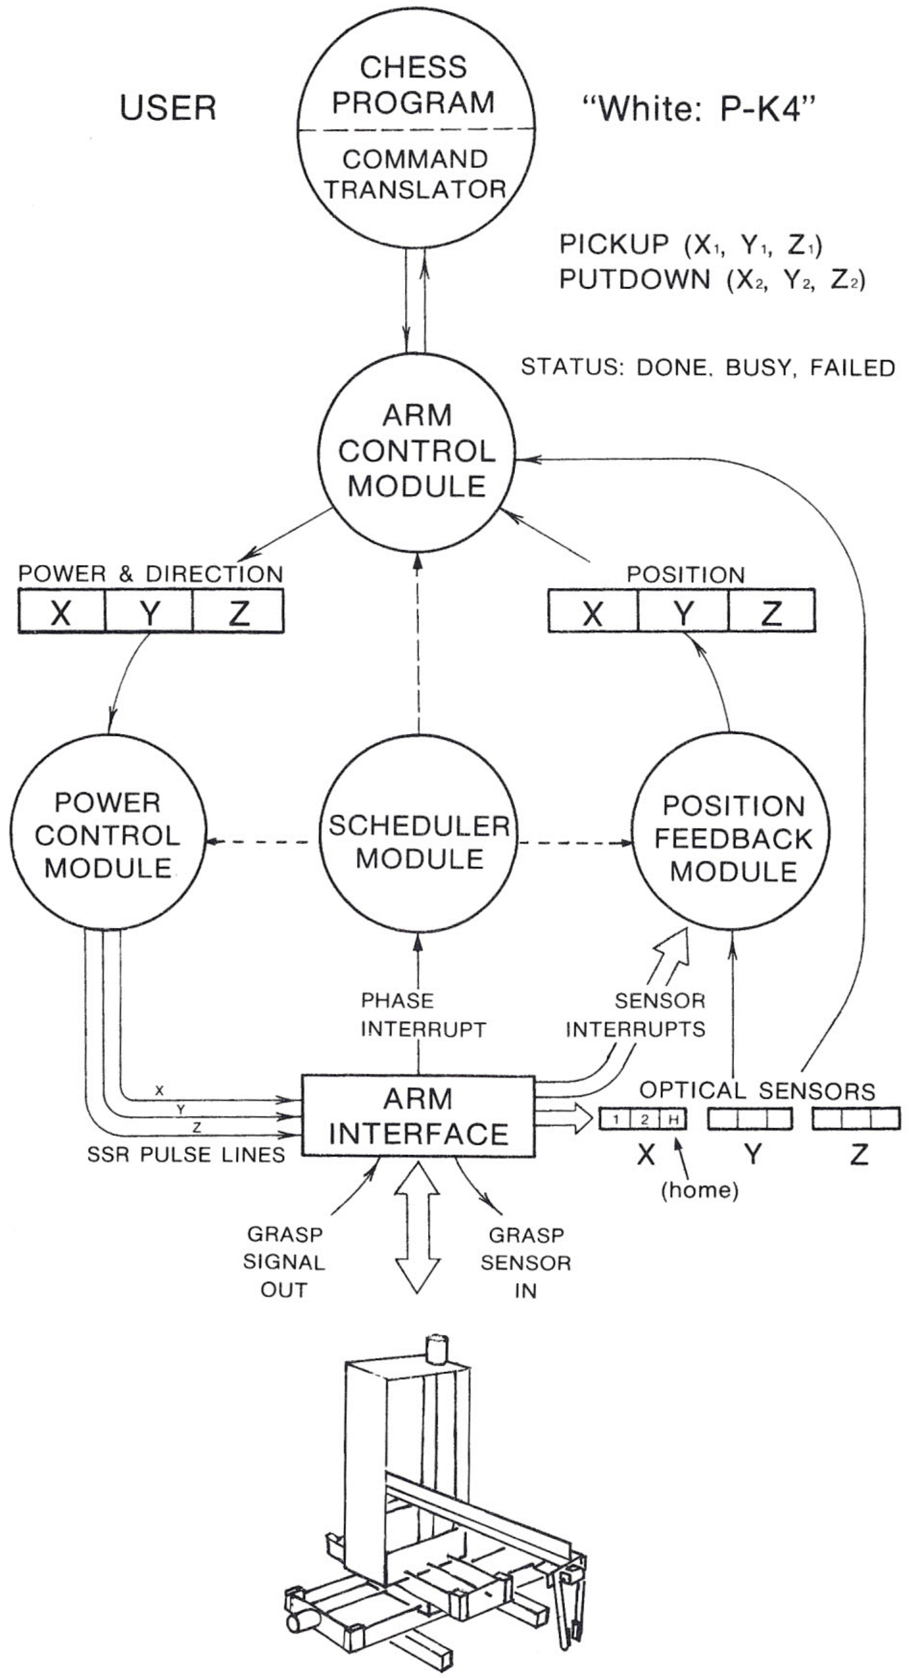

Figure 8. An interrupt-driven process configuration is used to control the arm. Modules communicate by shared control and feedback variables. The operation is coordinated by a scheduler module (center).

The power phase interrupt serves as a system clock whenever it occurs, the status of the interrupted module is saved and the scheduling module is run. Based on the priority (determined by their order in a process table) and status of the other modules, the scheduler decides which of the others should be run. This technique is common to most process control systems and time-sharing systems. In this system, the modules that are scheduled this way are the power control module, the arm control loop, and the chess program. In order of its priority, each module is checked to see if it is ready to execute. A module is ready to run if the "delay time" for the module is zero or if it had been interrupted by a higher priority process; if it is not ready, and its delay time is a positive number, then its delay time is decremented. After each module has been checked, the first one found ready is resumed. The execution of a module may terminate in one of three ways: it may reschedule itself (by setting its delay time) and return to the scheduler (in which case its status becomes "active but not ready"); or it may run until it is interrupted by the next clock cycle (remaining active and ready); or it may return to the scheduler without rescheduling itself, in which case its delay time is marked negative, and the module becomes inactive until some other process reschedules it.*

The arm control module operates as the decision function of a feedback servo loop that compares the coordinates given by the command translator with the current arm position and decides what the direction and power for each link should be. For example, if the arm is near the board and lateral movement is required to reach the target, then Z movement is needed to avoid hitting pieces on the board. If the arm had been receiving full power for long enough to reach appreciable speed, then when the arm gets sufficiently close to the target coordinates, the motor power should be reduced, or even reversed, to slow it down. If a link stops just a few turns away from the target, such as when it is seeking the home position, then the motor should be pulsed at very low power to close in without overshoot. If a grasp is indicated, it should close the fingers, delay for sufficient time for the fingers to close, check the grasp sensor, and search around if it missed. If it decides to change the power to a link, the arm control module sets the power command variable for that link and schedules the power control module to be run on the next clock cycle.

*Time-sharing systems use a different priority scheme that schedules each user according to the percentage of the CPU that he is entitled to at a given time.

The power control module examines the power commands (and other elapsed time counters) and decides which of the power circuits for the links should be pulsed at the time of the particular phase interrupt that caused it to be run. This module runs with greater priority than either the feedback loop or the chess program, since it is important to catch the power cycle at its beginning for a well-defined power pulse. (Using fractional cycles is an alternative control method, but, due to more precise timing requirements, is less suitable for this application.) After any necessary pulse commands have been sent to the arm interface, the module will reschedule itself to be run on the next clock cycle at which a pulse is needed.

The power control module can interrupt either of the lower priority processes. The position feedback module is activated by the sensor interrupts, and thus can interrupt any of the regularly scheduled processes regardless of priority. However, it must itself be present in the priority list since it can be interrupted by the clock and must therefore be resumed by the scheduler (unless it runs with the interrrupts disabled). Its priority is between that of the power control and arm control modules, but its delay time is set negative so that it will be ignored by the scheduler unless it had been interrupted. The chess program is the "background" task, getting any time left over by the other processes, so that even when the arm is moving it is still able to work on its search for the best move to counter each of the alternatives likely to be taken by its opponent. Of course, if the arm is not moving, it gets all the time, and can even turn off the "foreground" scheduler to reclaim the time that would otherwise be needlessly used to look for other active tasks.

This method of feedback control is referred to as a "discrete servo control loop" (as contrasted with the continuous servo loops implemented in analog circuitry), and the loop frequency needed to guarantee accurate control depends upon the response time (mechanical time constant) of the system. Given a particular combination of the inertia of the arm and the maximum torque of the motors used to drive it, the repetition rate of the arm control module can be gradually reduced until control of the arm begins to deteriorate (causing overshoot or oscillations).* The cycle time is determined by the delay

*The actual method of determining the proper feedback cycle time is complex, and includes choosing the proper way to use the feedback information (i.e., setting the position and velocity "gain" terms) as well as knowing the physical characterisitcs of the system. The references listed at the end of this article contain detailed discussions of this method of feedback control.

that the arm control module tells the scheduler to wait before running it again. The actual time between cycles is not precise, however, because the interruptions by the higher priority modules will be run first. The arm control module can estimate the speed of a link by counting the number of quarter turns that occurred since the last time it was run. This estimate can be used to determine the power to the link necessary to maintain a desired velocity or accelleration (using the control formulae discussed in the references).

An experimenter may want to add hardware to the system that performs some of the feedback or control functions. One choice could be to use digital up-down counters to count each quarter turn, maintaining the position register for the link and eliminating the position feedback software module and the sensor interrupts. By using a shaft position encoder of higher resolution and building a Phase-Locked Servo speed control circuit, the shaft speed can be controlled by hardware.* The GRIVET Series-3 arm kit can provide a testbed for exploring a variety of techniques both in software and in hardware. However, all the hardware needed for the basic control configuration discussed here is included in the kit, with the exception of the computer interface. The simplicity of the arm's design extends to its interfacing requirements — almost any parallel interface cards can be used, provided that the input board can be modified to produce the necessary interrupts. Writing the chess program, of course, is up to the ingenuity of the user! — Or one of the several microcomputer chess programs now on the market (Micro-Chess, etc.) can be modified for use with the GRIVET.

**See the article "Digital Control of DC Motors" in this issue of ROBOTICS AGE.*

REFERENCES:

— "Robotics — Grivet Series 1979 Catalog", Gallaher Research Inc., Winston-Salem, N.C.

On sampled servo design: (in order of increasing depth and technical level):

— J.A. Cadzon, *Discrete Time Systems,* Prentice Hall (1973).

— D.K. Frederick and A.B. Carlson, *Linear Systems In Communication and Control,* John Wiley & Sons (1971).

— B.C. Kuo, *Analysis and Synthesis of Sampled-Data Control Systems,* Prentice Hall (1963).

About the GRIVET Series

Some readers may remember seeing ads for the original GRIVET, which first appeared in March of 1978. This kit offered a wheeled chassis on which two manipulators were mounted. The arms were of the jointed variety—radial links connected by angular joints. Although the "Series-2" arm used in the original kit was discontinued at the end of last year, an improved version of the original wheeled chassis is still available. Gallaher Research Inc. will be offering the Series-3 manipulator discussed here for a limited time. Like its predecessor, the Series-3 arm was designed with experimenters, schools and small-scale industrial prototype evaluation in mind. Present plans are to offer this arm (and related hardware components in the 1979 catalog) until the end of the year, at which time the company plans to devote its resources to producing robot arms for light industrial application.

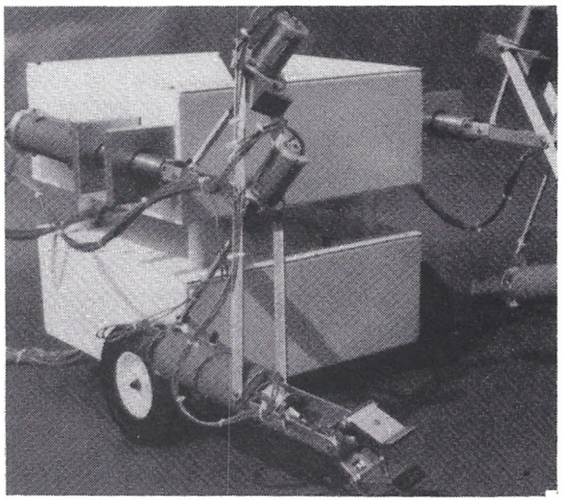

ROBOTICS IN THE SOVIET UNION

Many of our readers have inquired about the photograph of the Soviet hexapod robot which appeared in some of our promotional literature. Both this photo and that of the "Transport Robot" were hand-carried out of the Soviet Union by a visiting American computer scientist. We should point out, however, that there was really no intrigue involved; both robots have been described either in Soviet or international robotics and machine intelligence journals. But clear photos of Soviet robots are exceedingly rare, since copies of journal articles seldom offer good reproductions. News releases and translations of significant Soviet technical papers are often available to American scientists from the Federal Research Division of the Library of Congress, which is how we obtained part of our information on the hexapod robot.

The hexapod walking robot was developed at the Leningrad Institute of Aviation Instrument Design, and

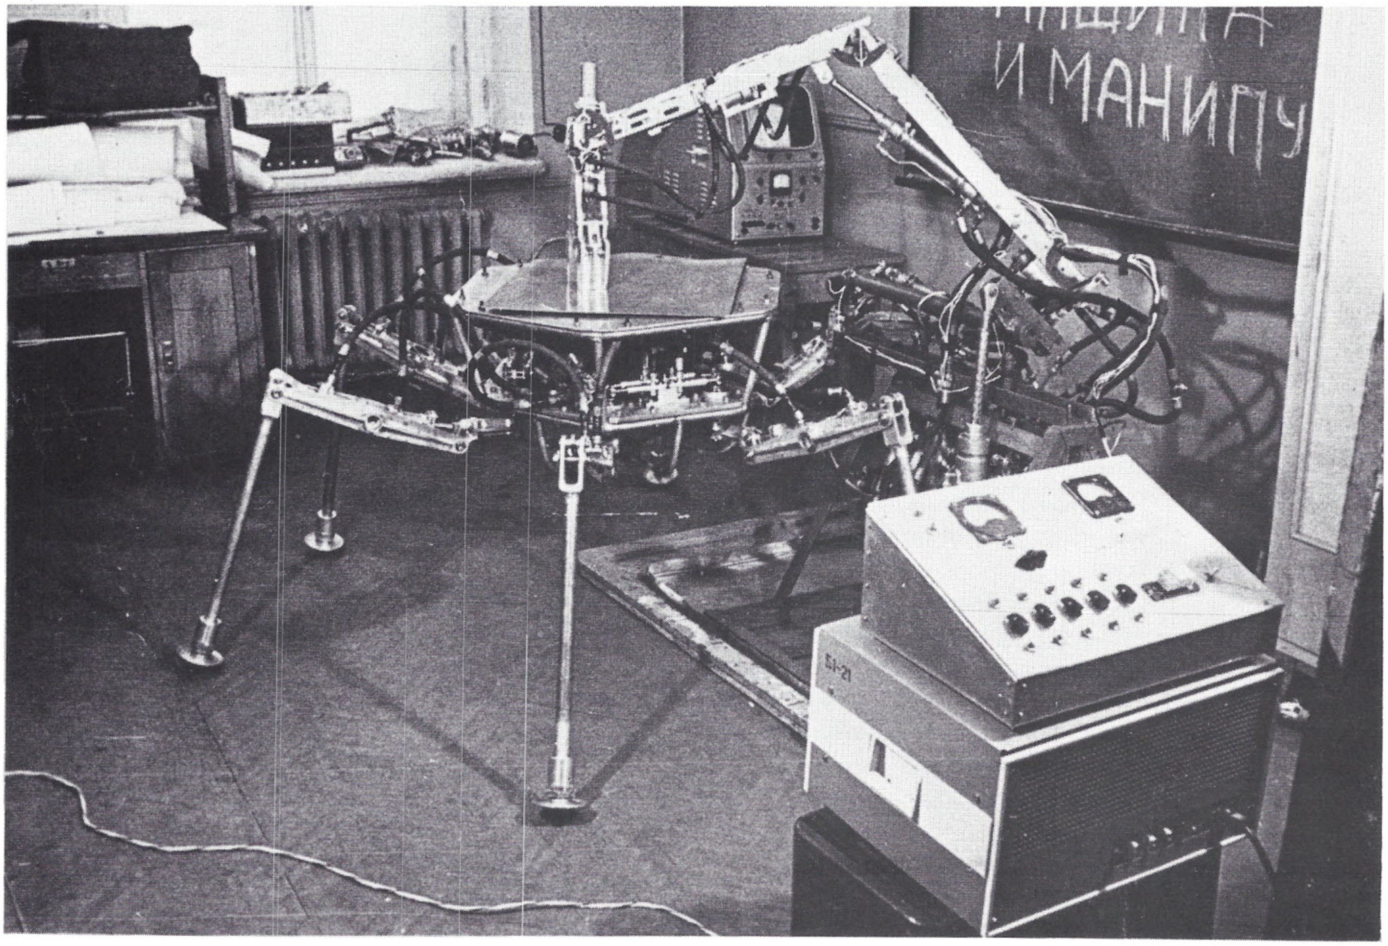

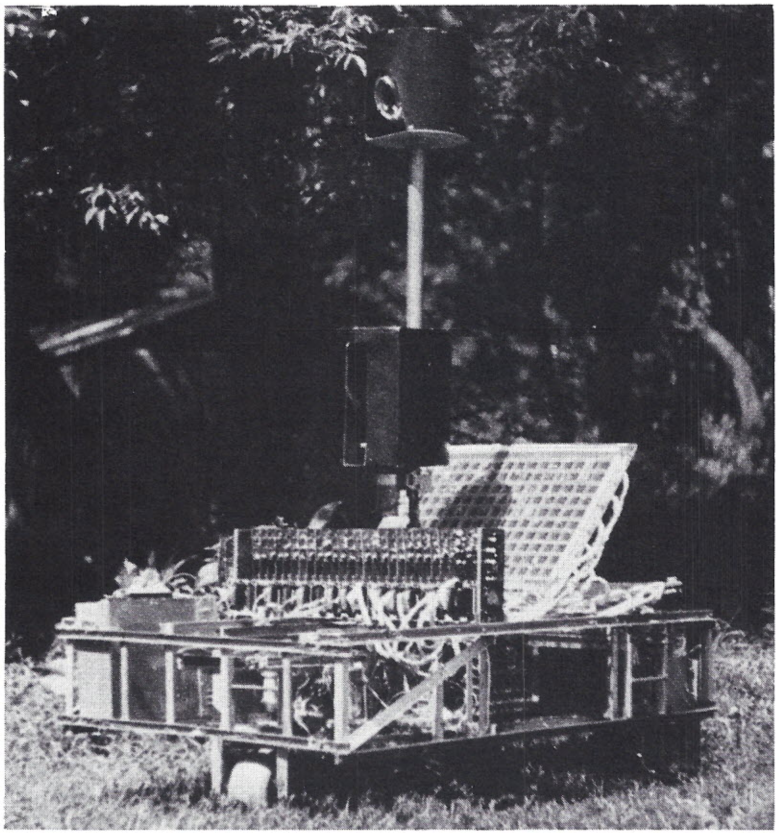

was first described publicly in the Soviet journal *Nauka i Zhizn'*. The robot is controlled by a central onboard computer, and is claimed to be capable of travel over almost any terrain by varying its gait and posture. The fastest gait of 6 km/hr is obtained by dividing the six legs into two stable tripods, and moving one tripod (and the body) forward while the other supports the body, then alternating. The joint angles can be varied to allow the robot to crouch or raise itself. A scanning laser rangefinder is used to acquire a local terrain map within a range of 10-15 meters. The scan rate and resolution may be varied depending upon the complexity of the terrain.

Technical papers describing Soviet research in walking robots were presented at the *Fourth International Joint Conference on Artificial Intelligence*, which was held in the Soviet Union for the first time in 1975 at Tblisi in the Georgia S.S.R. The Soviets are performing extensive research on walking systems, including two-legged robots. Some of the papers presented described techniques for choosing leg placement in rough terrains, and stability analysis of walking systems. The Soviet studies rely upon rigorous mathematical analysis of walking as a problem in control theory. Through such analysis, the position of the leg joint angles during a particular gait are determined by a control function in terms of a single parameter. The control rule includes the changes needed for the leg to accomodate variations in the terrain.

Soviet efforts to achieve "intelligent" control of robots rely heavily upon biological models. The "Transport Robot", developed by the Institute of Cybernetics of the Ukranian Academy of Sciences in Kiev, uses banks of "stimulus-response" cards (visible in the racks in the chassis) organized into synthetic "neural" networks. The robot, also described as IJCAI4, is trained by a reward/punishment scheme referred to as a "System for Reinforcement and Inhibition" (SRI), which alters the statistical "weights" of signals in the network.* The robot has a scanning photosensor mounted on a rotating mast which is used to detect obstacles in its path. The system has been trained to move toward a predetermined goal by avoiding obstacles.

Other Soviet robotics research is performed at the Bionics Laboratory of the Ukranian Academy of Sciences, the Institute of Applied Math, Moscow Ac. Sci., and the Siberian Ac. Sci.. Visitors to Russian laboratories report that the electronics in the research devices and even much of the test equipment are built by hand using discrete circuit elements. Integrated circuits and especially LSI devices are very scarce. Standardized function modules and boards (interfaces, etc.), common in the US, are practically nonexistent. Major research centers, however, have mainframe computers comparable to those in American universities.

* Similar systems were studied extensively in the US during the 1960's, but have had little application because of proven theoretical limitations and the potentials of other more general methods. See *Perceptrons*, by Marvin Minsky and Seymour Papert, The MIT Press, Cambridge, Mass., 1969.

ROBOTICS AGE COMPETITIVE EVENT

Progress in any field requires the efforts of creative individuals with the necessary resources to turn their ideas into reality. But creativity and resources in themselves are insufficient — the key to innovation is motivation. In most fields, the resources needed to pursue new goals are provided because of their profit potential. Sometimes, however, the goals may seem too far off to attract this investment. When this happens, extra stimulation in the form of contests and challenges can provide that extra push that can lead to the realization of those goals much sooner than would otherwise be possible. The races and cash prizes of the early days of the aviation and automobile industries are good examples of this tradition. Most recently, the challenge of the Kremer Prize led to the development of the first successful man-powered aircraft, the "Gossamer Condor," and the subsequent successful crossing of the English Channel by the same team in the "Gossamer Albatross." Consider also the enthusiasm and effort in the development of chess-playing computer programs stimulated by the Levy Challenge.

It would seem that the field of Robotics would not need such a contest to get it off the ground, but we think that there is a lot that a competition would contribute. Industrial and governmental investors in robots are looking for cost-effective solutions to numerous immediate problems. Industrial duty devices must be fast and reliable, and their large cost is justified by the greater cost of the alternatives. Also, the systems are typically designed for a particular application (although the availability of general-purpose assembly robots is growing). What we have not seen is a comparable development effort at the other end of the scale, namely the development of a general purpose domestic robot designed as a home appliance within the reach of most family budgets.

Certainly the goal of a "semi-intelligent" domestic robot with more than very limited sensory and manipulative abilities is a long way off. So long, in fact, that industry has little motivation for investing at this point, since the gradual development of industrial robots will inevitably lead to such capabilities anyway. We feel that Robotics will greatly benefit from a contest that will spark the imagination of individual inventors and set them to work on creating systems that can do useful work in and around the home by the innovative application of microcomputers and inexpensive electromechanical techniques. The field is ripe for this kind of initiative, since the prices of powerful micros and sensors have decreased to the point that meaningful experimentation is within the reach of student projects and many individual inventors.

For this reason, we at ROBOTICS AGE have decided to hold an annual ROBOTICS AGE Competitive Event, or **RACE**, the first of which will be held next Summer. Over the next few months we will define and standardize some objective robot performance criteria in a variety of categories. Some possible tests will be suggested below, but we also encourage our readers to send us their comments on the suggestions and their own recommendations. Details of the categories will be published in ROBOTICS AGE, and a complete rule package will be available for prospective entrants. Robots competing at the contest will be judged by a panel of experts who are professional researchers in Robotics, and cash prizes will be awarded to the winning systems. Depending on the response, we may need to hold regional eliminations.

Choosing the competition categories is not a simple task. We want to minimize the constraints on the environment so that the robots must sense and accomodate variations in the problem. Of course, the systems must be entirely self-contained, capable of performing in an electrically and optically isolated environment (although contest-approved power and

lighting will be available). This is to eliminate the possibility of remote control by humans. For the same reasons, command inputs to the robot will be suitably limited, in most cases allowing only program selection without human-supplied parameters. The tests will be designed to demonstrate the different aspects of robot behavior required to accomplish useful work in domestic applications. Robots need not compete in all categories, as separate prizes will be awarded for each. To make the contest more accessible to entrants with limited budgets, an attempt will be made to classify robot contestants according to cost, so that prizes can be awarded in different cost categories or ties can be resolved in favor of the cheapest system, etc. Similarly, the tasks can be structured with various degrees of difficulty so that an entrant could decide at what level to compete. Give us your suggestions about ways to establish a fair cost-related "handicap".

The first category is that of simple manipulative skills of the "pickup and place" variety. Robots must move objects from one location and place them in an orderly arrangement elsewhere. Sensory requirements will be minimized by requiring the robot to handle only one kind of object at a time, such as dishes, blocks, etc. For each trial, the robot will be informed in advance what kind of object it must handle, and the environment will be structured so that the object may be easily located by primitive sensors (feelers, reflectivity thresholds, etc.) The objects must be carried from one table to another and then, depending on the initial command, either stacked or lined up. Additional generality comes from the requirement that any of the given types of objects may be used in a trial. In this task, as in the others, the speed of performance will not be an issue, as long as the job is completed within a given reasonable time.

Another category will be for more complex manipulative skills. Robots in this category should be able to remove dishes from a countertop, place them properly in a dishwasher, and then clean the countertop with a brush or sponge, etc. A successful contender should also be able to wash windows using the same manipulator with a different tool. [Yes, we do windows!] Again, if necessary, sensory requirements will be limited by careful arrangement of the environment, i.e., objects to be picked up will be easily discriminated from the countertop and will not be occluded by other objects. Locating slots in the dishwasher holder may be difficult, but an exact specification of the holder will be given prior to the contest. For the window washing task, the location of the window may be given to the robot beforehand. The mechanism for maintaining proper contact between the window washing tool and the glass will be up to the ingenuity of the inventor. Scoring in both phases of the test will be based on the robot's degree of success with the task. (And you can be sure it will lose points for breaking dishes or windows!)

The next two suggested categories pertain to robot sensory ability in speech recognition and vision. Both of these areas are, in the general case, still unsolved research problems in Artificial Intelligence, so we expect that systems limited to using a robot's on-board computer(s) will be fairly limited. However, we do hope that you will surprise us! Speech understanding will be tested on a vocabulary of spoken numbers and commands dealing with sequences of simple movements. Apart from hearing, only contact sensing will be required, so that the robot should be able to execute commands such as "Turn left and so until contact." The words may be spoken by the trainer, but a normal speaking voice should be used. Scoring will be based on both the size of the working

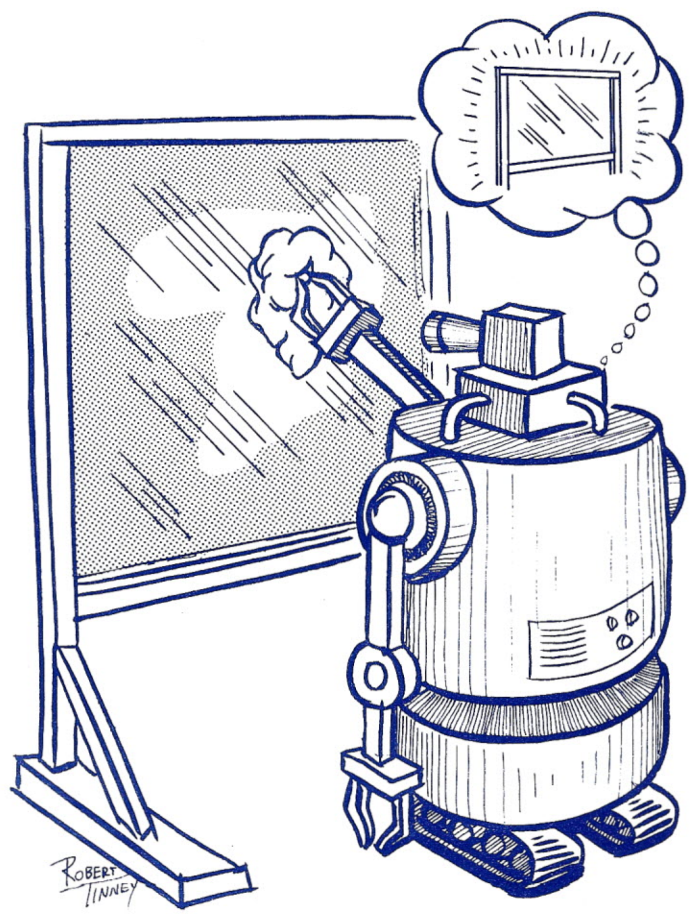

vocabulary and, to a lesser extent, the time required to recognize command words. The vocabulary and the permissible command syntax will be given beforehand, and the robot should be able to understand any legal sequence of words. The length of the pause between words should not overly influence the outcome, [Does anyone care to try recognizing continuous speech?] but any ties should be resolved in favor of the robot with the fewest CPU instruction cycles per word. Perhaps we should also require the robots in this category to have a speech generation capability as well. This would not add too much to the system cost, since several relatively inexpensive voice generators are commercially available (Votrax, Computalker, etc.) and would add a lot to the contest.

Vision is another difficult challenge for microcomputer implementation — most state of the art vision systems require either a mainframe computer or at least a large mini. Consequently, we expect most of the systems to be limited to two-dimensional image recognition.* For this category we will select a number of "standard" objects that the robots should be able to recognize, presumably from a known perspective defined beforehand. If several systems are successful at recognizing all objects under normal conditions (high contrast background, good lighting, etc.) elimination could be based on the robots' ability to recognize the same objects under worse conditions, i.e., less contrast, partial occlusion. If anyone wants to compete in three dimensional recognition (uncontrolled perspective and scaling) we will certainly make up a category if any other competitors can be found! The robot should indicate its selection of a particular object by moving or pointing towards it.

Perhaps the most exciting competition category is that of integrated robot systems. Robots in this category should have elements of all the other categories combined, although their sensory and manipulative capabilities need not be as advanced as in the robots competing in the specialized categories. They should be able to recognize a limited number of command words and be capable of visually detecting the presence of objects in high contrast conditions and discriminating between a few two-dimensional shapes. They should be capable of movement in an environment containing obstacles that must be detected and avoided, preferrably by passive (non-contact) sensing. Their manipulators should be able to reach from the floor to about a meter above it.

The tasks for this category should be varied, demonstrating the robot's flexibility. They would include such problems as finding clothes on the floor and picking them up, (extra points for hanging them up!) vacuuming around furniture or washing floors. The robot may be asked to go fetch a particular kind of object and bring it back, a task involving search, navigation, recognition and manipulation. If there is interest in competing in related outdoor tasks, we could arrange to have such problems as lawnmowing by visual navigation, harvesting fruit or vegetables, or washing cars. The robots should be basically general purpose devices that can perform different tasks by a change of program and by using different tools.

One of the reasons we're so excited about the competition is that none of the individual capabilities involved is at all beyond the state of the art in robotics. The challenge comes from having to bring the elements together in a microcomputer implementation contained in one package and to get it on wheels (or legs, if anyone is that ambitious!) Each capability, taken individually, requires no more processing power than an ordinary microcomputer with an adequate memory complement and appropriate device interfaces. Again, the challenge will be to design control systems that allow the robot to load the appropriate program when it is needed and to pass results from one software module to another. Manually changing floppy disks will probably not be allowed.

If the robot performs too slowly to suit you due to having to load programs in for each behavior function, consider using multiple CPU boards or Single Board Computers (SBC's) that are dedicated to particular functions and communicate by passing messages. A more ambitious system might use the new Zilog Z8000 chip set, which offers performance comparable to medium-sized minicomputers for only a few hundred dollars more than systems using its predecessors. This new generation of supermicros such as the Z8000 and Motorola's M68000 has one of the features of modern time-sharing systems that enables the system to automatically load in segments of program or data that are not in memory at the time they are needed.

Although research in visual image recognition has traditionally been confined to well-funded university and industrial laboratories, several developments during the last few years have combined to make TV input available to projects with modest budgets. To digitize a video input signal and to enable a computer to rapidly access the

**See the article "An Introduction to Robot Vision" in this issue. The General Motors 2-D recognition system described therein sets a standard for the sort of capability we would expect. To achieve that level of performance using an on-board microcomputer would be ambitious indeed!*

results used to require extensive design and construction of special digital hardware. Added to this was the expense of a TV camera. The growing use of video for security and home entertainment use has brought the prices of black & white TV cameras down to less than $250. The development of solid-state imaging arrays, now available from RCA, GE, and Fairchild Semiconductor, has provided devices highly suitable for robotics applications, due to the calibration accuracy possible with the silicon array and their compatibility with digitally clocked circuits. Home-built cameras using one of the smaller imaging arrays can provide excellent video input to your robot for a reasonable price.

As for getting the video signal converted to digital gray levels and into your computer, the options are not quite as varied. There are several commercially produced video digitizers that can digitize a single pixel or convert a TV frame and hold it for access by the CPU. [See the NEW PRODUCTS section.] Some allow simultaneous TV input and computer access for real-time applications. Video input to an S-100 bus machine can be purchased for as little as $350. Of course, you can beat these prices by designing your own TV input system, possibly using single pixel random access conversion, the "Slow-Scan" television that has become popular in Amateur Radio circles, or by using any of the high-speed A/D converters that are now on the market.

Voice recognition and manipulator design also present special challenges to the people who want to compete in those categories. Several voice input systems with various speed and vocabulary capabilities are on the market, including (at least) one with an S-100 interface which should do well in some of the **RACE** categories that require voice input. The next issue of ROBOTICS AGE will have a Robotics Product Index that will be of great use to those planning to build a **RACE** contestant. Products and manufacturers will be listed that relate to the design of each robot subsystem and capability involved in the contest. In addition, the articles in the coming issues of ROBOTICS AGE will be selected partly for their relevance to the contest. We have articles planned for voice recognition, power circuits, feedback control, manipulator design, and computer hardware/software organization in addition to articles on other aspects of robotics.

Now, you're probably wondering about the payoff! We have to admit that winning the **RACE** will probably not pay for your robot. We have industrial sponsors and other sources of revenue that will amount to approximately $10,000 in prize money. Bear in mind, though, that this will be divided among all the categories in the competition, so a lot depends upon the response we get from you as to how we design the contest. We can promise a lot of publicity, and industrial representatives will be present to look for new ideas. (We'll have a patent attorney standing by, just in case you're worried!) Be sure to send us your recommendations about the choice of the categories, the design of the tasks, and the means of judging. Remember that this will be your contest, because it is the individual inventor that we want to benefit from the encouragement that the contest offers.

We feel that the main thing that the **RACE** has to offer Robotics is a way of bringing together all the folks who have been working away in their basements, garages, and school labs with little communication with one another. We have seen the enthusiasm generated by recent microcomputer conventions and the explosive growth of the industry, and we hope to achieve the same result with the **RACE** and Robotics Convention next summer. See you there!

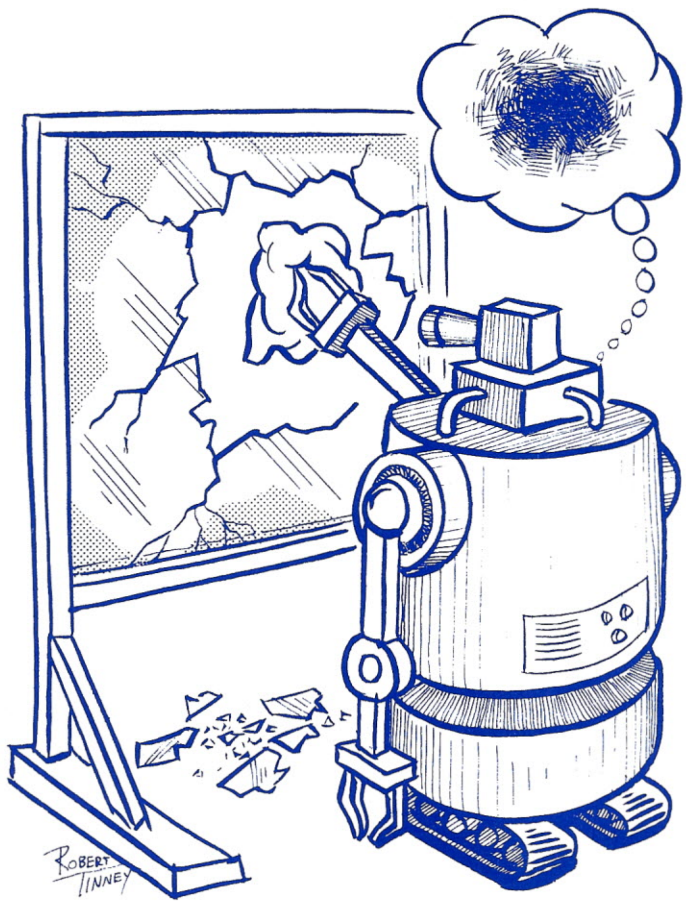

NEW PRODUCTS

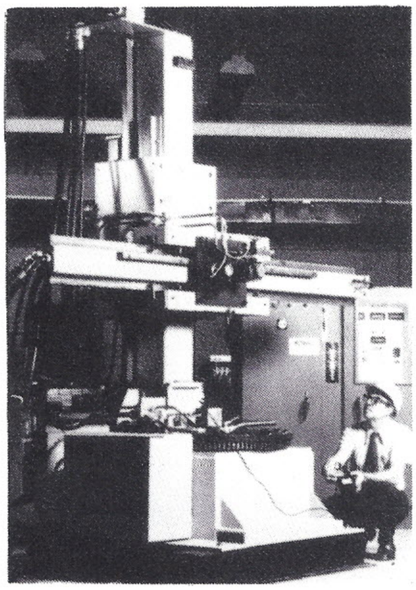

Robot Handles 2000 Pound Payloads

The Versatran FC Industrial Robot has weight handling capacity up to 2000 pounds. Manufactured by Prab Conveyors, Inc., this unit provides robot capabilities for a whole range of applications never before possible for industrial robots. Potential applications include foundry and machining applications where the combinations of part and gripper were too heavy to be handled by other programmable robots. In addition to these industries, many transfer line operations which were forced to use specially built automation for the transfer of heavy transmission cases, engines, and other automotive components can now opt for the advantages of a pre-engineered standard and programmable industrial robot.

The FC provides high speed straight line motions with smooth positioning. It handles bulky loads delicately with a high degree of accuracy, which is important when handling investment castings or parts with machined surfaces. The straight line (cylindrical) coordinate system used in all Versatran Models gives the capability of applying a 3 axis robot to applications where other robots often need 6 axis. This can eliminate the wrist motions which have traditionally been a source of problems when handling heavy loads.

The Versatran series of Robots offer from 2 to 7 axes of motion plus control systems ranging from simple cycling control to a 7-axis microprocessor-based controlled system. Both the Prab and Versatran Robot lines are manufactured by Prab Conveyors, Inc., Robot Division, 5944 E. Kilgore Rd., Kalamazoo, Michigan 49003.
Circle 1

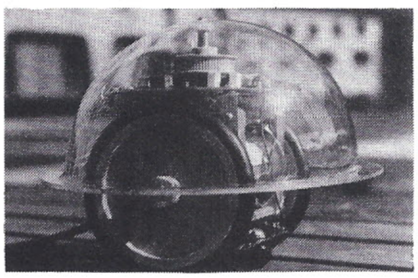

"Turtle" Robot for use with Home Computers

Cambridge, Mass.—R2D2 it's not, but the Turtle, a small home robot invented by William D. Hillis, a senior at MIT, may be the next best thing. The Terrapin Turtle is believed to be the first commercial offering of a home robot ready for sale. "It can 'walk,' 'talk,' 'blink,' 'feel,' and draw," says Dan Hillis. Hillis is affiliated with MIT's Artificial Intelligence Laboratory which has used types of "turtles" in various projects.

The Turtle acquired its name because it is approximately the size of a box turtle, consisting of a chassis topped by a hemispherical dome 3½ inches in radius. The Turtle has two large wheels and is capable of moving six inches per second. By turning the wheels in opposite directions, the Turtle can be made to rotate exactly on the spot. Other endearing features of the Turtle include beeping its speaker and flashing its lights.

More important, however, the Turtle can "feel" and draw. The former is accomplished by using its dome/shell as a touch sensor so that the Turtle "knows" when it bumps into an object. To draw, the Turtle lowers a pen from its chassis and moves. Under the control of a host computer (supplied by the user), the Turtle can map rooms by entering them, and storing the configuration in computer memory. Then the Turtle can let down its pen and draw a rough representation of the floor plans.

The Turtle kit retails for $300, from Terrapin, Inc., 33 Edinborough St., 6th Floor, Boston, Mass. 02111 617/ 482-1033. *Circle 2*

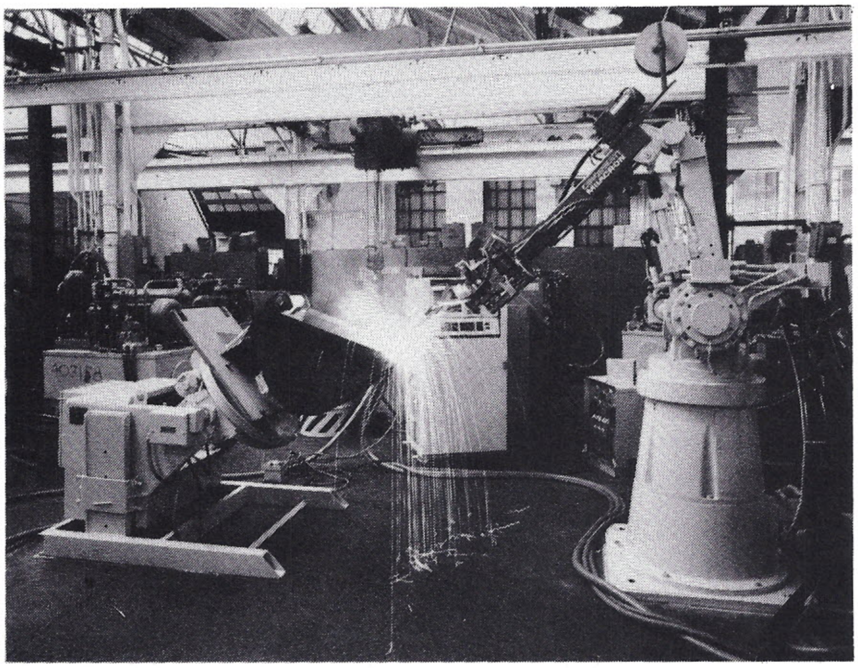

Milacron offers T³ Industrial Robot in "Turnkey" Automated Welding Packages

Cincinnati Milacron's computer-controlled T³ Industrial Robot is now being offered as part of a "turnkey" package for welding applications, tailored to various industries. When purchased this way, the package includes the robot plus all the ancillary welding equipment — power supply, welding gun, wire reels, feeder, part positioning table, etc. All necessary interfaces for optimum production are included.

The T³ is interfaced with Cincinnati Milacron's own minicomputer-based Acramatic CNC console. The robot is a fully articulated device capable of six axes of motion anywhere within a 1000-cubic foot work envelope surrounding it.

The robot is being offered with a new control configuration that takes up half the floor space of earlier versions while retaining all of the required control capability. The new control provides for storage of up to 700 different tool center points within a given program.

The T³ is simply programmed using a hand-held "teaching" pendant. The teaching function provides for three teaching modes — cylindrical, rectilinear, and hand-coordinated—wherein a single button on the teach pendant can control up to six axes of motion.

The software package provided by the new control features utility branching ability. This enables the operator to manipulate branch routines, through an offset feature, to alter the basic program in the control without actually reprogramming it.

Milacron has recently announced the availability of a new HT³ model, a heavy-duty version that can handle heavier loads with greater precision.

For more information, write Cincinnati Milacron, Cincinnati, OH 45209.

Circle 3

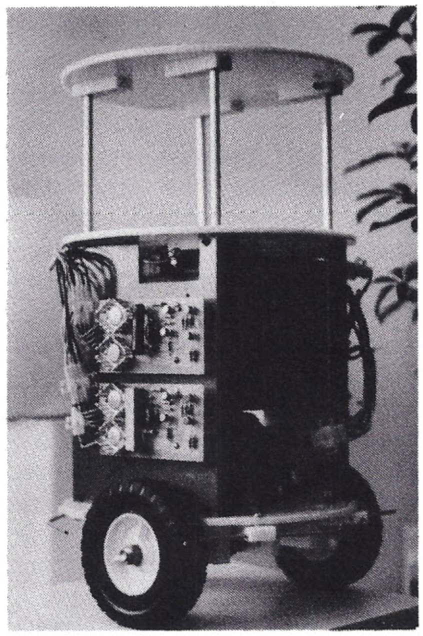

The Robot Shell Et-2

With the et-2 builder's manual from Lour Control you can build a rather sophisticated robot in your own workshop using widely available materials and off-the-shelf parts. et-2 (standing for experimental transmobile with 2 drive motors) can also be purchased in either kit or preassembled form. Made, for the most part, out of a rugged PVC plastic (other materials can also be used) et-2 is not a robot in the usual sense of the word, but a robot shell that requires the addition of some sort of control system to become a functional device.

The diametrically opposed drive wheels are powered by their own independent motor and drive train, while the rear of the shell is balanced by a single, heavy duty castor. The system is controlled by four command bits, which can be generated by standard TTL or CMOS IC's, though suitable buffers should be used, and are translated into motor response by the transistor switching circuits located on the two converter boards

NEW PRODUCTS

(one for each drive motor) along with the reversing relays held on the power/relay board. Those three cards make up the shell's interface system which, like the shell itself, employ no real exotic parts and can be easily fabricated using the schematics shown in the builder's manual. Complete boards, of course, are included with both the kit and assembled units. The relay/power board also includes six circuit protecting fuses and two power regulators that supply enough current to power the interface system as well as some modest controls. The shell also has a simple sensory system in the form of four mechanical bumper/switches that ring the lower tier. Each bumper is built around a dual-pole, dual throw switch that can be tied directly into your controls to trigger an avoidance response when the switch is activated.

Originally designed as a laboratory device, et-2 fits into the home as easily as it does into the lab. Standing just under two feet in height and about 14 inches wide, it travels with enough power to move easily on thick carpeting. To make et-2 more useful, Lour Control is developing a manipulator to mount on the front of the shell that will be able to pick up and carry small objects and even flip switches. Also on the drawing boards is an ultrasonic detector to supplement the mechanical bumpers the shell now has. The builder's manual, which leads you step by step through the shell's construction using both text and diagram, sells for $15.00 postage paid. More information on the kit and assembled units can also be obtained by writing the Lour Sales Group at 1822 Largo Court, Schaumberger, Ill. 60194. *Circle 4*

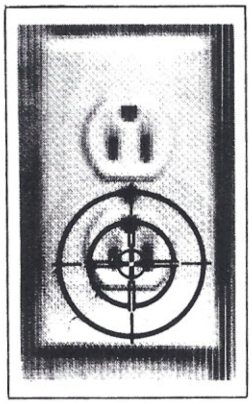

Random Access Video Digitizer

The MICRO WORKS is offering a new line of microcomputer-compatible video signal digitizers that allow computer software to selectively digitize any element of the picture. Since these digitizers are software driven, your system can extract only the information it needs to solve a problem—from detecting ambient light conditions to providing image input to pattern recognition programs.

The DS-80 (S-100 bus) and DS-68 (S-50 bus) Digisectors are random access video digitizers which provide the following features: a 256 x 256 picture element scan, 64 levels of grey scale, conversion times as low as 4 microseconds per pixel, accepts either interlaced (NTSC) or non-interlaced (industrial) video input.

Operation is simple; the computer sends the Digisector two 8 bit addresses (X and Y coordinates) and the Digisector returns the digitized brightness of the image at the specified location. Since digitization is one pixel at a time under software control, average conversion times are software dependent. The software supplied will digitize one pixel every other horizontal scan line, filling 8K of memory with two 4 bit grey scale values per byte in a little less than two seconds. The software is compatible with the Vector Graphic High Resolution Graphics Display Board. Additional routines are provided to drive both graphics and ASCII printers.

For setup and monitoring purposes the Digisector produces an output comprised of the camera's video signal plus a super-imposed cursor, showing exactly where the Digisector is looking. DS-80: $349.95 DS-68: $169.95. The Microworks, P.O. Box 1110, Del Mar, CA 92014 714-756-2687. *Circle 5*

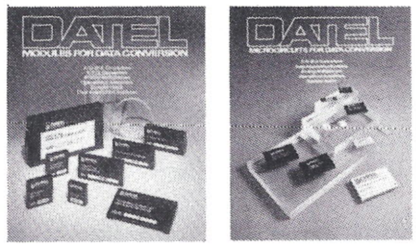

Two Free Catalogs by Datel Systems

"MICROCIRCUITS FOR DATA CONVERSION", a 48-page catalog by Datel Systems describes in detail a broad line of high performance monolithic and hybrid data converters. A companion catalog, "MODULES FOR DATA CONVERSION," describes data conversion modules and accessory circuits. The product line encompasses many new state-of-the-art products from A/D and D/A converters to analog multiplexers, sample-holds, fast operational amplifiers, voltage to frequency converters and active filters. These devices are specifically designed for a wide range of measurement, control and instrumentation

applications. The second catalog also has a section devoted to the principles of data acquisition and conversion.

Datel System's monolithic and hybrid devices are available in three different operating temperature ranges: 0 to +70°C, -25 to +85°C and -55 to +100°C or +125°C. A stringent quality assurance program has been implemented, and high reliability devices screened to MIL-STD-883B are available. Datel Systems, Inc., 1020 Turnpike St., Canton, Mass. 02021. 617/828-8000

Circle 6

Electromechanical Linear Actuator

Clifton Precision, Special Devices, Litton Systems, Inc., 5100 State Road, Drexel Hill PA 19026, has expanded its product line of Linear/ Rotary Actuators and Servo Systems for military and commercial applications.

Model 7501-03 Linear Actuator, designed to operate the Ram Air Door on wide body aircraft, features continuous stall, high reliability, low backlash and high-force/weight ratio. The unit can also be used in a closed-loop control system. Other rotary/linear actuator systems are available up to 5HP for use in aircraft, helicopters, missiles, drones and RPV's.

The company also produces Rare Earth/Alnico Permanent Magnet DC and brushless motors, stepper motors and 400 Hz servomotors; film and wire-wound potentiometers; absolute and incremental encoders; and servo electronics. For further information, contact Russ Philipp, Marketing Product Manager, 215/622-1000 Ext. 396.

Circle 7

CAT-100 digitizes a video image in 1/60th of a second. Shown is an image displayed as 227 lines by 288 pixels of 4 bits.

S-100 Plug-in Unit Digitizes TV input, Displays Digital TV Images in Color

The first complete video imaging and graphics system with real-time TV frame grabber which plugs into the S-100 bus has been introduced by Digital Video Systems of Palo Alto, CA.

The CAT-100 is a compact two-board color video imaging system providing three fundamental functions: a video frame digitizer, an image memory, and an output video generator. The digitizer can capture a video frame in 1/60th of a second and store it in the on-board 32K-byte image memory. The video generator displays the digitized image in 16 shades of gray or 16 colors on standard low-cost black and white or color TV monitors. It can also directly drive a high-resolution RGB monitor.

The CAT-100 can generate its own RS-170 video synchronization, or automatically extract it from any external video source such as a TV camera, video disc, video recorder, live off the air, or from a separately supplied external synchronization signal.

Two types of video A/D conversion circuits provide a choice of 1, 2 or 4 bits per pixel at maximum video rate. The first circuit is a conventional converter which yields 16 gray levels. The other converter is a contouring circuit useful to reduce images to binary outlines by automatic thresholding.

The software-selectable system parameters include a variety of image formats, video output controls, digitization commands,

NEW PRODUCTS

addressing modes, vertical image offset, photographic trigger and lightpen control. Users can choose from 15 full screen formats, including square, rectangular and condensed aspect ratios. These formats are available for digitization as well as for display, an resolutions range 256 to 1280 pixels per TV line. Typical formats include 240 lines of 256 pixels of 4 bits, and 480 lines of 512 single-bit pixels.

The 32K-byte image memory is fully and permanently accessible for image generation or processing in the address space of the S-100 bus through a 2K-byte window which can be dynamically selected or deselected by software. This window can be located on any 2K boundary in address space.

The CAT-100 also provides an excellent tool for text editing and processing. Six different formats are available in a text mode as an alternative to the graphics and imaging mode. The characters are generated on a high definition 7 x 9 matrix, and up to 2,640 characters can be displayed on the screen, organized as 33 lines of 80 characters. This densely packed display only represents 8 percent of the available 32K-byte memory. A smooth scrolling feature allows the user to scan the entire text file at a variable speed while the characters remain perfectly legible. Any portion of the 32K-character text file can be displayed by setting the appropriate offset address in an instruction register. Any number of cursors can be defined by the user.

A lightpen can be connected directly to the CAT-100. The interaction is extremely precise and the 18 bits of coordinates provided by the CAT-100 actually resolve one pixel in the 480 x 512 format. The lightpen allows the user to easily create or modify digital images.

The photographic trigger input is another original feature that allows one to take clean individual pictures of the screen and to generate black and white or color animated films. The user's software library includes fundamental driver routines, lightpen tracking, image enhancement, mathematical point and line plotting, and a versatile software character generator with controllable parameters.

The CAT-100 can be upgraded by adding memory extension boards. A buffer capacity of 64K bytes or more yields 8-bit pixels, i.e. 256 levels or colors per pixel.

Applications include image enhancement and analysis, automated counting and measurement, quality control monitoring, character and pattern recognition, education, medical and scientific research, surveillance, simulation, computer portraiture, architectural design, TV studio effects.

Nine different options are available. Prices range rom $893 to $1992.80 according to configuration. A typical unit with a 4-bit frame grabber and 32K-byte image memory sells for $1222 (1). Quantity discounts. Delivery: stock to 3 weeks. DIGITAL VIDEO SYSTEMS, 595 Matadero Ave., Palo Alto, CA 94306. 415/494-6088, J. Robert Flexer. *Circle 8*

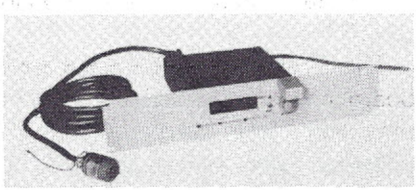

Scalable Absolute Encoders for Severe Environments

A series of low-cost, scalable absolute encoders especially suited to severe electrical (EMI) and physical (high temperature, vibration and moisture) environments has been introduced by Computer Conversions Corp. These new units will convert any shaft input to BCD or binary information corresponding directly to the shaft angle with an accuracy of ±1 pt. in 3600. In addition, 3, 4 or 5 digit, .5"H LED displays of that angle are available. The shaft transducer used is a highly reliable resolver.

Multiple shafts can be built simultaneously encoded at a very low cost per channel, simply by adding another transducer and part of the single channel electronics package. Data outputs are TTL compatable and a data transfer and data hold line are provided for simple interfacing to a computer. The basic update rate is 2.5 ms.

Multiturn units also are available. Pricing for Pt. #DS90-DB-10-2 is $375.00/axis in unit quantities. Delivery - 4 weeks ARO. CCC is located at 6 Dunton Court, East Northport, New York 11731. Phone: (516) 261-3300. *Circle 9*

Video Digitizer gives Computers "Eyes"

A Microprocessor-based video image analyzer, plug compatible with Data General computers, is being introduced by Octek Inc. of Burlington, Massachusetts.

The Octek Model 2000 is a microprocessor based, video image analyzer that can simultaneously digitize and store (on-board) video from a TV camera or other image source in real time, enter it into any host Data General computer, and display it on a TV monitor. The self-testing, 15" x 15" card also performs graphics generation, image

processing, and archival image storage (when used with a disc memory). If used with a modem, it can transmit TV images over voice grade telephone lines.

With a host computer access time of 800 ns per pixel, the Octek Model 2000 offers a resolution of 320 pixels (X-axis) by 240 pixels (Y-axis). To emphasize or de-emphasize image features, the 16-level gray scale is modifiable by an input/output computer programmable video lookup table. Some typical applications include automatic inspection, surveillance, product sorting, traffic control, OCR, video animation, map reading, and area and volume analysis.

The Octek Model 2000 is unit priced at $4700 each, including RDOS software driver; OEM discounts are available. Literature and technical data will be provided on request. Contact: John E. Trombly, OCTEK Inc., 121 Middlesex Turnpike, Burlington, MASS 01803, (617) 273-0851.

Circle 10

Personnel Consultant Servicing the Robotics Industry

The President of Uptrend Associates, Inc., Mr. Frank M. Dyer, announces that a Robotics Industry Division will be added to their engineering, sales, and data processing search divisions. This new division will work exclusively with client companies engaged in the manufacture and sale of industrial robots. Mr. Dyer states that the formation of this division was prompted by their recognition of the specialized requirements of the robotics industry.

"All our projections," Mr. Dyer reports, "indicate that the demand for personnel with robotics expertise will increase at a rapidly accelerated rate. Companies desiring to obtain and hold the lead edge in robotics technology will be in constant search of highly specialized and scarce personnel." Uptrend Associates, Inc., has over ten years experience in search and placement for firms engaged in electronics, computer, instrumentation and other areas involving the design, manufacture and sales of high technology mechanical and electromechanical equipment. Uptrend Associates, Inc., 4120 Southwest Frwy. Ste. 114, Houston, TX 77027, 713/623-2020.

Circle 11

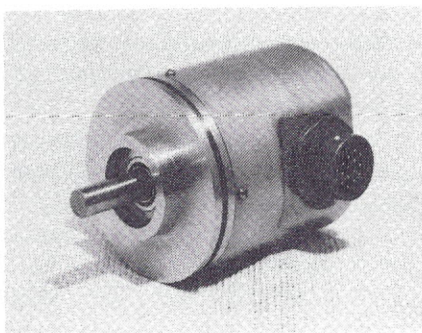

Absolute Code, Optical Encoder Model 76

Litton Encoder Division, Chatsworth, CA has introduced a new line of low cost, absolute code, shaft position encoders.

The encoders are supplied in gray code, natural binary, or 8421 binary coded decimal. Model 76 offers a choice of ten resolutions, with DTL and TTL compatible outputs. The rugged frame is available in three mounting configurations. The units are compact; measuring 4 inches by 2.65 inches in diameter. The solid-state illumination source is guaranteed for five years.

Typical applications for Model 76 include NC machine tools, computing scales, material handling systems, antennae, navigation systems, and a variety of other uses. Additional options assure this unit the adaptability for varying customer needs. Litton Encoder Division, Marketing Department, 20745 Nordhoff St., Chatsworth, CA 91311, Phone: 213/341-6161.

Circle 12

New Book Analyzes Hearing

The Bionic Ear, by Dr. John L. Stewart, PhD., offers a thorough analysis of biological and synthetic hearing systems. The book is primarily intended for design engineers, physicists, artificial intelligence researchers, and audiologists.

Part one, Analysis, studies the structure and function of the ear using differential equations, decibel measure and frequency response functions, statistical detection theory and the Weber law.

Part two, discusses optimal procedures for hearing. The kinds of signals required for an animal to perceive for survival during major evolutionary periods are estimated. Design of an optimum system to detect and classify these stimuli is undertaken.

Part three concerns ear models starting with the classical form of the differential equations and simple ladder network representations of the cochlea, and applies concepts in network theory and analysis. Inadequacies of classical models are discussed, and active circuit representations are introduced which better describe the nonlinear ear, with designs for practical systems.

Part four describes actual and potential applications having

NEW PRODUCTS

considerable economic significance, and means for reducing speech to the simplest set of measures are described. Critical measures that result relate to phoneme identification, marking phoneme boundaries, and pitch inflections. Practical results include band-width compression, aids to computer speech recognition, and aids for the handicapped that could permit sound recognition by tactile stimulation. Practical computer voice input devices are discussed.

The author is a former researcher in the Air Force Bionics Program, and has written three textbooks and numerous journal papers in related work. The book is available for $35, from Covox Company, Box 2342, Santa Maria, Ca., 93456 805/937-9545.

Circle 13

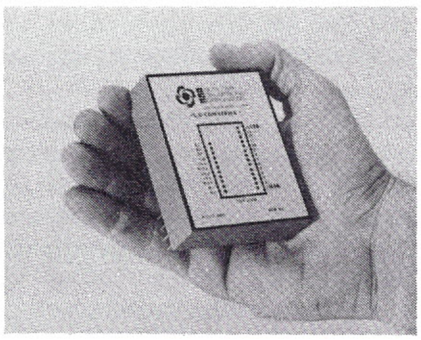

Synchro to Digital Converter With DC Velocity Outputs

A series of Synchro to Digital Converter Modules which provide DC voltage output proportional to velocity in addition to a 10, 12 or 14 bit digital output have been introduced by Computer Conversions Corp.

The ±10V DC velocity outputs can be used as feedback signals in closed loop servo systems and also to replace tachometers. The new modules are 2.6″ x 3.1″ x 0.82″ and are designed to be mounted on P.C. boards. They convert synchro or resolver inputs of 11.8v or 90v, 400Hz or 90v, 60Hz., into parallel binary outputs representing angle with an accuracy of up to ±4 minutes of arc. There is no accuracy degradation over the operating temperature range, ±10% amplitude and frequency variations and ±5% power supply variations. Input rates up to 10,000° per second can be tracked with no added error.

The converters have isolated reference and synchro inputs and provide complete synchronization to a computer. The digital outputs are DTL/TTL compatable and bi-directional input data is accepted. Low power schottky logic units are also available.

Prices in reasonable production quantities are less than $350.00 each. Delivery is 4 weeks ARO.

CCC is located at 6 Dunton Court, East Northport, New York 11731 (516) 261-3300 TWX #510 2260448.

Circle 14

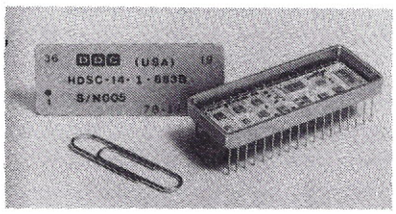

Industry's Smallest 10-bit Converter is Hybrid Design

The HSDC-10 is announced by the ILC Data Device Corp. as the industry's smallest synchro-to-digital/resolver-to-digital tracking converter. The entire hybrid unit is contained in a single double DIP package.

The rugged converters, which are processed to MIL-STD-883 Level C are also optionally available to Level B standards.

The new converter series at 400 Hz has a Ka of 200,000 $sec.^2$ minimum, and the settling time to 1 LSB for a 179^0 step is 25 ms typical. Standard input options include all the normal three-wire inputs and four-wire resolver formats. The unit accepts broadband inputs 360 to 1000 Hz, and features high CMR. Output angle is normal binary code, parallel positive logic and TTL/DTL and CMOS compatible.

These new DDC converters are accurate to ±21 minutes, or ±1 LSB in 10 bits. This accuracy, which includes quantizing error, is maintained under all static and dynamic conditions at speeds up to ±100 rps at 400 Hz or 20 rps at 60 Hz. Using CMOS logic and low-power IC's results in a power consumption of less than 200 ms. In test, MTBF's are as high as 1.6 million hours.

Measuring only 0.78 by 1.9 by 0.21 inches, the 18-gram HSDC-10 is factory adjusted and tested over the full range of its operating temperature (55° to 125°C, or 0° to +70°C). These new converters require no field adjustment or calibration.

A 14 bit version (HDSC-14) is available as two 36-pin double DIP packages (Ka = 58,000, power is 750 mW, ±4′ accuracy, ±0.9 LSB worst case error, with control transformer and error processor in separate packages.)

A companion 14-bit hybrid D/S converter converts digital angle to either synchro or resolver output format with ±4 minute or ±2 minute accuracy. Contained in a single 36 pin DDIP module, the HDSC-14 is not only the smallest D/S or D/R

converter available, but is also a versatile function generator. It can convert digital angle to sine/cosine DC, generate a cartwheel rotating sweep for PPI displays, or convert from polar to rectangular coordinates. Other features include a short-circuit proof ±2 mA rms output which is pin programmable for synchro or resolver mode, and high CMR. The digital input is transient protected CMOS with internal pull-up resistors. The digital input is CMOS and TTL compatible. Output settles to final value in less than 20 microseconds for a 180° step.

HDSC-14 prices start at $495 for quantities of 1 to 9 (U.S. and Canada) and delivery is stock to 8 weeks. ILC Data Device Corporation, Airport International Plaza, Bohemia, New York 11716 516/567-5600. *Circle 15*

Datel Enters Switching Power Supply Market

Datel's entry into the switching power supply market is marked by the introduction of its first model, a 25-watt line operated module. Model USM-5/5 has a 5 VDC regulated output at 5 amperes and operates at 80% minimum efficiency; the supply is encapsulated in a standard 2.5 x 3.5 x 1.25 inch compact module.

The design of the USM-5/5 incorporates the use of monolithic switching regulator· and Schottky rectifiers. Using a push-pull pulse width modulated technique which operates at 20 KHz minimum results in silent operation.

Other features include a wide input line voltage range of 90 to 130 VAC at 47 to 450 Hz. Output ripple is 50 mV peak to peak maximum. Output voltage is 5V 1% with overvoltage protection and current limiting short circuit protection.

Line regulation is 0.05% and load regulation is 0.1%. Output voltage temperature coefficient is 0.02%/°C maximum. The operating temperature range is –25°C to +71°C with the output derated 2%/°C from 42°C to 71°C. Warm-up drift is .02% at full load. The USM-5/5 operates up to 42°C ambient temperature at full load with normal convection cooling.

Applications include powering TTL logic circuits, semiconductor memories and data acquisition and conversion systems. The high efficiency permits generating main system power at +5 volts and then using a DC-DC converter from the +5V line to produce ± 15 VDC for analog circuitry. Price: (1-9) $109.00 Delivery: 4-8 Weeks ARO. Datel, 1020 Turnpike St., Canton, Mass. 02021, 617/828-8000. *Circle 16*

Relay Multiplexed A/D Board

Datel Systems introduces the ST-711RLY A/D board series with a relay multiplexer specifically for industrial environments with noisy signal leads, and high common mode voltages.

This ST-711RLY series is mechanically, electrically and program compatible to Intel's SBC-series one-board microcomputers, and gets its name partly from full register and pinout compatibility to the popular SBC-711/732 and ST-711/732 series fast A/D-D/A boards. Applications include slowly-varying parameters such as temperature, pressure, flow rate, etc., using low level bridge-type transducer-sensors such as thermocouples, strain gauges, load cells, etc. For these applications, the 30 sample-per-second throughput is entirely adequate, and the selectable ranges from 10 mV to 2 Volts are ideal. Datel uses the computer's ±12 VDC power bus to run the A/D section. This eliminates the DC/DC power converter normally found on A/D boards. Also included are board pads for current shunts, overload clamps or attenuators. A/D conversion is to 12 bit binary resolutions with accuracy of .03% FSR (2V range), and .1% FSR (10 mV range). CMR is 126 dB, from 0 to 60 Hz with 2.5 pF and 5 x 10^{13} ohms to ground. System throughput is 36 mS, nonlinearity is ± ½ LSB, differential input impedance (leakage) is 20 megohms and will accept ±15V overvoltage (sustained). For real-time applications, a jumper-selected Pacer Clock timebase is included to externally start A/D scans or interrupts. The board size is 12"W x 6.75"D x 0.5"H (305 x 171 x 13mm) and weighs 22 oz. (0.6kg). The operating range is 0 to +70°C. The ST-711RLY8D (8 channel) is $650 in singles and the ST-711RLY16D (16 channel) is $995 in singles. Availability is 8 to 10 weeks ARO. Datel Systems, Inc., 11 Cabot Boulevard, Mansfield, MA 02048, Telephone (617) 828-8000 or (617) 339-9341. *Circle 17*

New Products Policy

ROBOTICS AGE encourages the manufacturers of products related to robotics and automatic control applications to send in New Product announcements for this department. Readers are advised that, although ROBOTICS AGE selects new products for publication based on our estimation of their relatedness to the field, in most cases we have not evaluated either the product or the company and, although we would not knowingly print any inaccurate information, no endorsement of any product is implied. Send announcements of New Products literature to ROBOTICS AGE, New Products, P. O. Box 4029, Houston, TX 77210.

TECHNICAL ABSTRACTS

As part of our goal of disseminating current technical information to our readers, this department will list abstracts of significant recent technical papers, in cases where these papers are available to the public. Some research institutions do not ordinarily distribute their papers outside of the professional research community, and we will attempt to arrange for ROBOTICS AGE to offer reprints of robotics-related papers we believe will be of interest to our readers. We strongly suggest that, out of consideration for the author's time, all correspondence regarding these documents be directed to the publications departments of the institutions involved. The relevant addresses will normally be listed after the abstracts. We urge academic and industrial research centers to send us abstracts of recent papers in Robotics and Artificial Intelligence for possible inclusion in this department, with appropriate prices and ordering procedures.

A Modular Vision System for Sensor-Controlled Manipulator and Inspection, by Gerald J. Gleason and Gerald J. Agin, SRI TN-178 ($3.00)

This paper describes a prototype hardware/software system for industrial computer vision applications. The hardware consists of a solid-state TV camera, a preprocessor interface, and an LSI-11 microcomputer. The software consists of routines for analyzing binary images including the following: connectivity analysis, extraction of shape descriptors, automatic recognition, training for recognition by showing, and determination of position and orientation. The system has computer-controlled dual thresholds, which are useful in generating brightness histograms. Actual processing time for a typical image is reported.

Real Time Control of a Robot With Mobile Camera, by Gerald J. Agin, SRI TN-179 ($3.00)

In this paper we describe and analyze a control system for a Unimate robot that derives its control information from a small, solid-state TV camera attached to its end-effector. Visual input may be in either of two modes: using conventional lighting or using projected degrees of freedom that may be controlled and their attainable accuracy image, the resolution of the camera, and the relationships between the camera, target, and projector (if used). A dynamic analysis of the system accounts for discrete delays in the control loop as well as the transfer function of the robot itself. The system has been demonstrated in several modes simulating manufacturing operations in static and moving coordinate systems.

An Interactive Incremental Compiler for More Productive Programming of Computer-Controlled Industrial Robots and Flexible Automation Systems, by William T. Park and David J. Burnett, SRI TN-180 ($3.00)

SRI International, in conjunction with Philips of the Netherlands, has developed a general-purpose software package for fast interactive programming of computer-controlled flexible automation systems, such as industrial robots. This software could materially reduce the set-up time for batch production runs.

The system is written in the RTL/2 programming language and runs on a Philips P857 32K-work, 16-bit minicomputer, with Philips' DOM disc operating system. Robot programs written in an RTL/2 subset are partially compiled into interpretable object code, then interpreted.

The interpreter supports single-step operation, execution tracing, breakpoints, and data value displays. The programmer-trainer can interrupt a running program to change data values and revise the algorithm. If no active statements are changed, the robot can then continue working without having to start over. Eliminating restarts allows the programmer-trainer to spend more time maximizing the robot's production rate, because there is no need to replace the workpieces, tools, and manipulator arms in their starting position every time a small change is made in the robot's program.

Hierarchical Representation of

Three-Dimensional Objects Using Verbal Models, by Gerald J. Agin, SRI TN-182 ($3.00)

We present a formalism for the computer representation of three-dimensional shapes, that has as its goal to facilitate man-machine communication using verbal, graphic, and visual means. With this method, pieces may be assembled hierarchically using any of several ways of specifying attachment. The primitives of the representation are generalized cylinders, and the creating of assemblies may make use of the axes inherent in the primitives. Generic models may be described that may leave some parameters or dimensions unspecified, so that when a specific instance of the model is described, those parameters may either be explicitly specified or take on default values. The axes of local coordinate frames may be given symbolic names. A set of computer programs translate descriptions of objects into polyhedral models and line drawings.

Robotic Sensors In Programmable Automation, by David Nitzan, SRI TN-183 ($4.00)

There exists both a social and an economic need for the advancement of automation in general, and of programmable industrial automation in particular. Characterized by flexibility and the ease of setup for new production tasks, programmable automation employs industrial robots. Today's robots possess "muscles" only; there is a need to develop intelligent robots that can detect faults and correct errors by using sensors and computer control. Intelligent robots may have contact sensors (sensing force/torque, touch, position, etc.), noncontact sensors (sensing visual images, proximity, range, etc.), or both. Such sensors are applicable to three basic functions: inspection, finding objects, and robot-control feedback. Fast reaction to multi-sensory data can be achieved by using a distributed network of microcomputers that process the sensory data in parallel and control each effector in a modular fashion. Application of sensor-controlled manipulation to material handling, inspection, and assembly tasks has been demonstrated in several laboratories in the United States, Japan, and Europe.

Minicomputer Software Organization for Control of Industrial Robots, by William T. Park, SRI TN-184 ($3.00)

Two mainstreams of robot software organization are evident in the arm-control programs written in recent years. In the first, programs model the arm's working environment in a large, sophisticated computer. This achieves a certain level of understanding of the work so that people do not have to give such detailed orders. In the second, programs require more explicit instructions from a person, but they can run in a much smaller and less expensive computer.

SRI Technical Notes may be obtained by sending a list of the desired TN numbers, together with your payment (prices are listed in parentheses after the TN no.) to:

Georgia Navarro
Technical Note Librarian, J2057
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

LETTERS

Dear Editor:

Congratulations on being the first to recognize the need for a journal that specifically addresses the rapidly growing field of robotics.

John E. Trombly
OCTEK, Inc.
V.P., Marketing

Dear Editor:

As an avid roboticist, I found the ad for your magazine fascinating. I am definitely interested in subscribing, and so I enclose a check.

Sincerely,
Douglas Zimmerman
Finksburg, MD

Dear Editor:

Just wanted to take time out to wish you luck on your first issue of ROBOTICS AGE.

Gene Beley

Dear Editor,

ROBOTICS AGE seems like it might be a rewarding wealth of information, so therefore I would like to start my subscription. I have been involved with robotics on the hobby level for over 22 years and it is now becoming a full-time business.

Sincerely,
Tom Carroll

(cont. pg. 66)

MEDIA SENSORS

NEWSWEEK, April 23, 1979 - **BLUE COLLAR ROBOTS** Bernard Sallot, executive director of Robot Institute of America, says robots are good workers because they take no coffee breaks, never get pregnant, never go out on strike, draw no pensions, and uncomplainingly do various nasty jobs that a human worker would disdain. Less than two decades old, industrial robotics is now a $60 million-a-year business and could grow to $3 billion-a-year by 1990.

Joseph F. Engelberger, president of Unimation, Inc., says that in 1961 one of his $25,000 robots cost about $4.20 an hour on the job, while human labor was $3.80. Today the robot costs $4.80 an hour while the human costs about $14 an hour.

General Motors Corp. has about 150 robots welding auto bodies at its Lordstown, Ohio assembly plant. By 1990, Engelberger predicts, robots will have displaced about 5% of the blue-collar work force that now performs "the least desirable of jobs." Engelberger also says that $50 million in research & development, plus a high final unit cost, could give us robots for household chores by 1990.

(AP) Nashville, Tennessee **'HERMAN' THE 'ROBOT'** The Department of Energy's Oak Ridge, Tenn. office has a seeing (twin t.v. cameras) remote controlled robot that runs on "tank-like" treads, manipulates objects with its arm, and is attached by wires to a control trailer. It was flown to the site of the Three Mile Island nuclear power plant accident. Used to enter areas of lethal radioactivity, Herman was only needed for some simple radiation count and visual scan work at the plant. But he was once used in Rochester, N.Y. to remove a piece of radioactive cobalt stuck in a pipe, and once again in Sewanee, Tenn. when a radioactive source fell out of a container at the University of the South and had to be rebottled.

WALL STREET JOURNAL **$6.8 MILLION ROBOT PLANT** Cincinnati Milacron (machine tools, plastic goods, industrial chemicals) is planning a large plant to make parts for its line of industrial robots. James A.D. Geier, president, says "Our robot sales more than doubled in 1978 and we expect sales to triple this year...the new plant is necessary to meet the growing demand."

WALL ST.JOURNAL, April 16, 1979 **NEW CIVIL SERVANTS** Mail Robots in Federal Offices - 20 mail robots prowl the corridors of 10 federal agencies in D.C. today, saving taxpayers money and bureaucrat's time. Robert Heinemann, a Commerce Dept. administrative offficial, says, "They're terrific. They don't go to lunch and never go to the bathroom." January 1977 was the first time a mail robot began working for the government, delivering mail over a one mile long floor route. Mr. Heinemann says, "They proved so beneficial that we bought five more." All are called Robby. The Mailmobiles are produced by a division of Lear-Siegler Corp. It looks like a five-foot-high cart with pigeonholes and blue lights. But four six-volt batteries make the 700-pound robot move, and invisible lines of chemicals painted on the floor keep it on track as it moves along a route of drop-off points, pausing 20 seconds and chiming at each.

The Air Force tried a 'Norman the Mailer' unit in late 1977, but ceased for fear of security leaks. Says an Air Force spokesman, "...you just don't have the human element watching over the classfied stuff."

But at the Labor Department, two more units are being purchased. And at Commerce, one 'Robby' has been modified with $150 worth of stereo equipment to play music while making his rounds.

Lear-Siegler has sold over 300 of its mail deliverers to private companies. They even tested a Mailmobile that is capable of wandering down a hall sweeping, brushing and polishing the floor. But the company is concentrating on the potentially larger market they see in mail robots.

WALL STREET JOURNAL, May 2, 1979 - **ROBOTS MAY BECOME CRIME FIGHTERS** According to George Wackenhut, president of Wackenhut Corp., a major security and investigative concern, "Devices once associated only with the world of science fiction" are being studied by his company, to fight "crime in all its devious forms." Some of the

company's new tools may include "intelligent androids" that can be used for security duty.

HOUSTON POST, March 19, 1979
NEWEST GENERATION OF ROBOTS Robots are being used more rapidly in high-wage countries like America, Japan, Australia & Continental Europe. The robots sometimes recover their capital cost in less than three years. They do the dirty, dangerous jobs humans don't want. They don't mind working night shifts and can be "retrained" within minutes. They do tasks accurately, repeatedly, without tiring or complaining, and without absenteeism.

However, one school of thought says that existing machine tools will evolve - acquiring manipulative skills and incorporating sense of touch, sight and possibly hearing. Such a tool could listen to itself, and diagnose when it needed repair.

A step towards such advanced machine tools is a press with a built-in robot for loading and unloading sheet metal. The next step might be a robot that is being developed at Nottingham University. It manipulates workpieces under a hammer, dispensing with the costly dies used in pressing today.

The biggest need for robots is in large factories. Materials spend 95% of their time in factories being pushed from place to place, or just lying idle. One of the most advanced industrial robot systems to speed production is 'Robo-gate' at Fiat's Rivolta and Cassino factories. This machine welds a whole car body in one minute. More than 50 robots are combined in a welding bay. They not only replace costly multi-welding machines; but also allow an entirely different car model to be made on the same line, without expensive and time-consuming retooling.

BOOKS

How to Build a Computer-Controlled Robot, by Tod Loofbourrow, (Hayden Books, Rochelle Park, New Jersey) is a step-by-step instruction manual on how to build a microcomputer controlled robot capable of sensing its environment and responding under program control. It is a *true* robot — both the power supply and the micro are on board. It is designed to accomplish the most with the least financial investment, and with expansion in mind. You can build your "Microtron" (or "Mike", as Tod calls it) up to any of the three levels of increasing capability detailed in the book, and continue to add to it for as long as your ingenuity (and finances!) hold out.

Mike's controller, a MOS Technology KIM-1 microcomputer, is responsible for all motor drive signals and sensory processing. Stage I of construction consists of building the frame, installing the motors, drive circuitry, and computer, and connecting a joystick for testing. When completed, the Stage I robot will be able to follow preprogrammed courses with full internal control of speed.

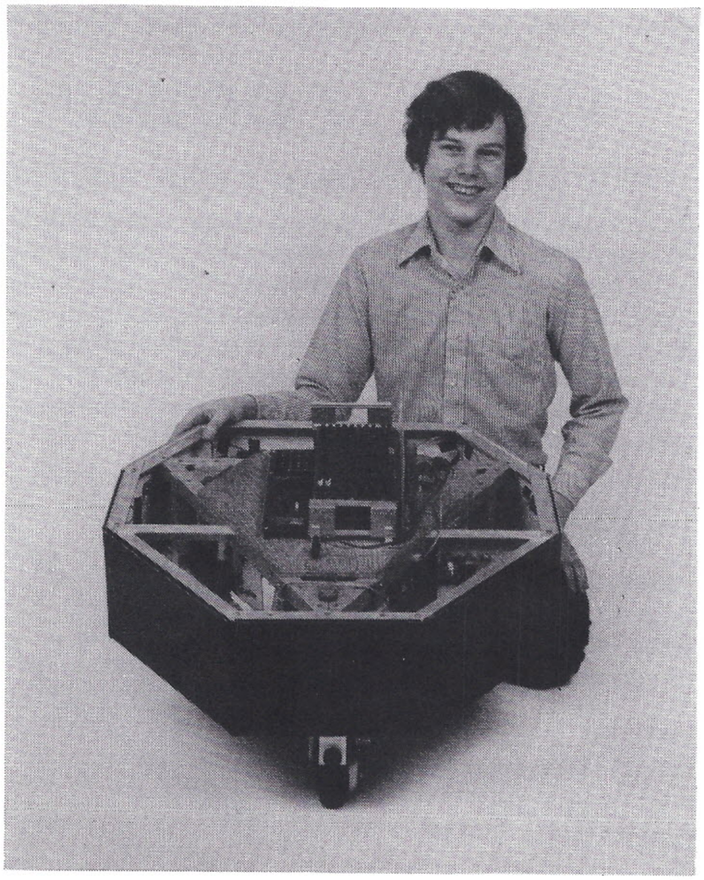

The next stage gives Mike ultrasonic sonar and contact sensors. Using inexpensive, readily available transducers, the sonar allows the robot to "see" objects up to ten feet ahead and to automatically divert to avoid collisions. The outer frame is surrounded by contact sensors that cause an interrupt upon collision, immediately initiating a sequence of commands which can cause Mike to back away.

In Stage III a simple but effective sound recognition system is added, giving Mike the ability to recognize up to eight different words using a program supplied by the author. Zero-crossings of the sound's waveform in both high and low frequency bands are counted in a time interval to produce a numeric "signature". In "teach" mode these signatures are averaged and stored in a table, so that in normal operation Mike can compare the signature of a sensed sound or word to what was previously taught, select the closest match, and initiate the appropriate response associated with the command.

Throughout the book we were

impressed with the clarity of the instructions and the quality of the drawings and schematics. Even more impressive, however, are the clever circuit and software designs Loofbourrow uses to accomplish sensory and control fuctions in the KIM with only a few inexpensive components. Many builders will want to experiment with more elaborate methods, and the Microtron provides an excellent testbed for this, but the basic functions are all provided in the designs — together with the KIM-1 code to use them. Sources of all the less common hardware components are listed. Building your own "Mike" from these plans should be very rewarding, and we are interested in hearing about the experiences our readers have with the project.

The Complete Handbook of Robotics by Edward L. Safford, Jr. (Tab Books) "Incidentally, Kryton, the Chief-Emperor-King Robot of them all, tells me math is a must subject if you really want to become an electron brother to the extended-life-form-robot, and android."..."We will have a chapter on the Black Art of Servo-Mechanisms."..."Now with Kryton and Klatu and all the gnomes and elves, and trolls and wizards and magicians standing by, being ever critical of everything we do, we begin the examination, study, and evaluation of this wonderful and delightful subject—robots."

These three quotes from the preface characterize the technical quality of *The Complete Handbook of Robotics*. For the most part, this book is just the sort of pseudo-technical jibberish that the field of robotics and anyone interested in it can well do without.

To be absolutely fair, there are probably a few dozen pages of genuine content, including definitions of numerous robotics terms and very superficial discussions of robot systems and techniques, even a few schematics (with no discussion) of circuits marginally related to robotics. The real problem with the book is that these few pages of content are spread, a line or two at a time, over some 358 pages of practically context-free text and mostly irrelevant photographs.

All things considered, you would probably be better off waiting until Reader's Digest comes out with a condensed version. Even if they decide not to pick it up and you miss it altogether, well, nothing lost. There is certainly a need for a book that honestly attempts to give the reader a broad, introductory survey of the field, but unfortunately this book is not the one.

The Psychology of Computer Vision, Patrick H. Winston, ed., McGraw-Hill, New York (1975).

The transfer of technical results from university research laboratories to the general public is a difficult process. Each year thousands of Doctoral theses and journal articles are published, but seldom do they reach outside the academic community. Textbook publishers must select for wider dissemination those few results that, over a period of a few years, have come to be regarded as significant milestones in the development of a particular branch of scientific knowledge. As a result, even though the material published in a hardcover technical book may be somewhat dated, it usually describes techniques whose validity have been proven by subsequent research. Such is the case with the early steps toward computer vision made at the MIT Artificial Intelligence Laboratory, using high-contrast scenes of children's building blocks as the subjects of analysis.

The "Blocks World" provided the opportunity to study the problem of analyzing scenes from an arbitrary three-dimensional perspective. Constraining the domain to scenes of blocks permitted the development of basic analytic techniques effective in the simplified domain, without having to face the complexity of scenes from natural environments at the outset. Although the work has received some criticism for its lack of generality, fundamental mechanisms were indeed discovered that suggested directions for further research and in a sense provide a standard by which more general systems can be measured. This book contains reprints of five significant technical papers, two describing blocks world techniques, two discussing techniques suggested by the blocks world, but with general applicability, and a theoretical paper on methods of organizing visual (and other) knowledge.

The work by David Waltz illustrates how structure observed in one part of a scene can be used to constrain the interpretation of another part, thereby reducing the complexity of the task of recognition. The system, operating on line drawings of scenes containing building blocks in arbitrary arrangements, searches for plausible interpretations of the vertices and edges of blocks in the scene. Any vertex considered individually may have many possible

BOOKS

interpretations — it may be on the inside or the outside of the object, the "matter" of the object may lie on one side or the other, etc. When any two vertices are joined by a common edge, the possible interpretations of the pair are greatly reduced. Given all the pairs of vertices connected by common lines, the labelling process rapidly converges, usually to a unique interpretation of the scene. The system also is capable of interpreting the apparent edges caused by shadows in the scene.

In the second paper, Yoshiaki Shirai presents a system in which an edge detection function is guided by knowledge about the kinds of objects that the scene may contain, in this case polyhedral blocks. After first filtering the image with a conservative edge detector, the system tries to find edges in areas of less contrast by varying the sensitivity of the detector. The program attempts to extend boundary lines into the interior of the image of a pile of blocks, to try to locate the boundaries between stacked blocks. Also, once the corner of a block is located, it attempts to find other edges that terminate at the corner. The result is a line drawing of the scene.

Berthold Horn provides a thorough mathematical treatment of surface reflectivity, showing that, if both the reflectivity function of the surface and the position of the light source are known, the shape of the surface can be derived. Such a capability underlies the human ability to use shading as a depth cue and has other scientific (and cosmetic!) applications.

A system developed by Patrick Winston demonstrates the ability to learn the meaning of structural terms from examples carefully chosen by a "teacher." Starting with an initial vocabulary of shape categories, structural and spatial relational terms, the system is able to extend this vocabulary by analyzing examples of new structures. The process involves comparing the structural description of a given example with the stored description of a similar structure. The differences between the two descriptions may then be used as distinguishing features to define different kinds of structures. Although the examples are taken from the blocks world, the method applies to symbolic descriptions in general.

Marvin Minsky's paper, "A Framework for Representing Knowledge", discusses a theory of organizing an associative relational memory. A "frame" can be thought of as a data structure that contains "slots" for features and facts normally associated with a particular situation. One advantage of this form of organization is that slots may have "default" values that define a stereotyped instance of the frame. The theory has proved useful in several subsequent computer systems.

This book will be useful for anyone interested in Artificial Intelligence, either for its historical value or as a means of gaining insight into both the methodolgy of AI research and the kinds of processes that underly robot vision. It is not a computer "cookbook," vision is too rich a field to attempt that in a single volume with any but superficial treatment, but the 3-D blocks world is described in detail. A possible criticism is that only the work done at MIT is presented, but again, no single volume should attempt to be encyclopedic. Appropriate references to related work and a computer vision bibliography are supplied.

LETTERS

(cont. from pg. 61)

Dear Editor:

I just received a brochure describing your magazine and would like to investigate the possibility of stocking ROBOTICS AGE on a regular basis in our store.

We have noticed a good deal of interest in robotics among our customers. While [other computer magazines] have occasional articles related to the field of robotics, and we do carry several books on the subject, I think there is definitely room for a magazine devoted to robotics. Please send information on pricing, terms, and quantities for establishing a dealership for ROBOTICS AGE.

David Hughes
COMPUTER SHOPPE
3225 Danny Pk.
Metairie, LA 70002

COMING SOON

IN ROBOTICS AGE

- Coverage of the 6th International Joint Conference on Artificial Intelligence -- Tokyo, 1979
- An Economical Voice Input System
- The Ćuk High-Performance Switched-Mode Servo Amplifier
- The Rensselaer Polytechnic Institute's Experimental "Mars Rover"
- A Low-Cost Video Input System
- Software Techniques for Microprocessor Servo Control

and MORE!!!

PRODUCT NO.	PRODUCT
1	Versatran FC Industrial Robot, p.52
2	"Turtle" Robot Kit, p.52
3	Milacron T^3 Industrial Robot, p.53
4	Et-2 Robot Shell, p.53
5	Random Access Video Digitizer, p.54
6	Datel Systems Catalogs, P.54
7	Electromechanical Linear Actuator, p.55
8	CAT S-100 Plug-in Unit, p.55
9	CCC Scalable Absolute Encoders, p.56
10	Octek Video Digitizers, p.56
11	Uptrend Assoc., Inc., personnel consultants, p.57
12	Litton Optical Encoder Model 76, p.57
13	*The Bionic Ear,* by Dr. John Stewart, PhD., p.57
14	CCC Synchro to Digital Converters, p.58
15	Hybrid Design HSDC-10 Converter, p.58
16	Datel Switching Power Supplies, p.59
17	Relay Multiplexed A/D Board, p.59
18	Uptrend Assoc., Inc., personnel consultants, p.63
19	Posters by R. Tinney, p.45
20	Chess Poster by R. Tinney, p.35
21	CAT-100 16 Color Imaging System

ARTICLE NO.	ARTICLE
1	Digital Speed Control of DC Motors, p.4
2	Industrial Robotics '79, p.12
3	Introduction to Robot Vision, p.22
4	The Grivet Chess-Playing Arm, p.36
5	Robotics in the Soviet Union, p.48

READER SERVICE

PLEASE SEND INFORMATION ON ITEMS CIRCLED BELOW. REQUESTS CANNOT BE HONORED UNLESS ZIP CODE IS GIVEN. PLEASE PRINT. THIS CARD IS VALID FOR 90 DAYS ONLY.

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30

NAME

TITLE

COMPANY

ADDRESS

CITY

STATE ZIP COUNTRY

READERS ADDRESS THE EDITOR

THIS IS YOUR CHANNEL TO THE EDITOR. MORE ARTICLES WILL BE INCLUDED IN FUTURE ISSUES BY AUTHORS AND ON SUBJECTS WHICH RECEIVE HIGH RATINGS.

ARTICLE NO.	POOR		GOOD		EXCELLENT	
1	0	1	2	3	4	5
2	0	1	2	3	4	5
3	0	1	2	3	4	5
4	0	1	2	3	4	5
5	0	1	2	3	4	5

ROBOTICS AGE
P.O. BOX 4029
HOUSTON, TEXAS 77210

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY CARD
FIRST CLASS PERMIT NO. 17539 HOUSTON, TEXAS

POSTAGE WILL BE PAID BY ADDRESSEE

ROBOTICS AGE
P. O. BOX 4029
HOUSTON, TEXAS 77210